Quantum

CAPACITY FOR THE EXTRAORDINARY™

Quantum Bigfoot™ CY 2.1/4.3/6.4 AT Product Manual

November 22, 1996

81-113668-01

Quantum reserves the right to make changes and improvements to its products, without incurring any obligation to incorporate such changes or improvements into units previously sold or shipped.

You can request Quantum publications from your Quantum Sales Representative or order them directly from Quantum.

Publication Number: 81-113668-01

UL/CSA/TUV

UL standard 1954 recognition granted under File No. E78016

CSA standard C22.2-950 certification granted under File No. LR49896

TUV Rheinland EN 60 950

Tested to FCC Rules for Radiated and Conducted Emissions, Part 15, Sub Part J, for Class-B Equipment.

SERVICE CENTERS

Quantum Service Center
715 Sycamore Avenue
Milpitas, California
Phone: (408) 894-4000
Fax: (408) 894-3218
http://www.quantum.com

Quantum Asia-Pacific Pte. Ltd.
50 Tagore Lane #b1-04
Singapore, 2678
Phone: (65) 450-9333
Fax: (65) 452-2544

PATENTS

These products are covered by or licensed under one or more of the following U.S. Patents: 4,419,701; 4, 538,193 4,625,109; 4,639,798; 4,647,769; 4,647,997; 4,661,696; 4,669,004; 4,675,652; 4,703,176; 4,730,321; 4,772,974; 4,783,705; 4,819,153; 4,882,671; 4,920,442; 4,920,434; 4,982,296; 5,005,089; 5,027,241; 5,031,061; 5,084,791; 5,119,254; 5,160,865; 5,170,229; 5,177,771; Other U.S. and Foreign Patents Pending.

Copyright© 1996 by Quantum Corporation. All rights reserved. Printed in U.S.A.

This product or document is protected by copyright and distributed under licences restricting its use, copying, distributing, and decompilation. No part of this product or document may be reproduced in any form by any means without prior written authorization of Quantum and its licensors, if any.

RESTRICTED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraphs (c)(1)(ii) of the Rights in Technical Data and Computer Software clause at DFARS 252.227-7013 and FAR 52.227-19.

THIS PUBLICATION IS PROVIDED "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTIULAR PURPOSE, OR NON-INFRINGEMENT.

Quantum Bigfoot CY® is a trade mark of Quantum Corporation and AIRLOCK, WriteCache, DisCache, and the Quantum logo are registered trademarks of Quantum Corporation.

PC-AT is a registered trademark of International Business Machine.
MS-DOS is a registered trademark of Microsoft Corporation.

Table of Contents

Chapter 1
ABOUT THIS MANUAL 1-1
1.1 AUDIENCE DEFINITION 1-1
1.2 MANUAL ORGANIZATION 1-1
1.3 TERMINOLOGY AND CONVENTIONS 1-1
1.4 REFERENCES 1-3

Chapter 2
GENERAL DESCRIPTION 2-1
2.1 PRODUCT OVERVIEW 2-1
2.2 KEY FEATURES 2-1
2.3 STANDARDS AND REGULATIONS (Pending) 2-3

Chapter 3
INSTALLATION 3-1
3.1 SPACE REQUIREMENTS 3-1
3.2 UNPACKING INSTRUCTIONS 3-2
3.3 JUMPER SETTINGS 3-3
3.3.1 Cable Select/Slave Present (CS/SP) Jumper 3-5
3.3.2 Drive Select (DS) Jumper 3-5
3.3.3 Park Jumper 3-5
3.4 IDE-BUS ADAPTER 3-6
3.4.1 40-Pin IDE-Bus Connector 3-6
3.4.2 Adapter Board 3-6
3.5 MOUNTING 3-7
3.5.1 Orientation 3-7
3.5.2 Mounting Screw Clearance 3-8
3.5.3 Ventilation 3-9
3.6 COMBINATION CONNECTOR 3-9
3.6.1 DC Power 3-10
3.6.2 IDE-Bus Interface Connector 3-10
3.7 DRIVE INSTALLATION 3-11
3.7.1 For Systems With A Motherboard IDE Adapter 3-11
3.7.2 For Systems With An IDE Adapter Board 3-11
3.8 TECHNIQUES IN DRIVE CONFIGURATION 3-13
3.8.1 Maximum Partition Size in MS-DOS 4.0 and Above 3-13
3.8.2 1024 Cylinder Limitation on Older Computer Systems 3-13
3.8.3 Newer Computer Systems with Extended BIOS Translation 3-13
3.9 SYSTEM STARTUP AND OPERATION 3-14

Chapter 4
SPECIFICATIONS 4-1
4.1 SPECIFICATION SUMMARY 4-1
4.2 FORMATTED CAPACITY 4-2
4.3 DATA TRANSFER RATES 4-2
4.4 TIMING SPECIFICATIONS 4-3
4.5 POWER 4-4
4.5.1 Power Sequencing 4-4
4.5.2 Power Reset Limits 4-4
4.5.3 Power Requirements 4-5
4.6 ACOUSTICS 4-6
4.7 MECHANICAL DIMENSIONS 4-6
4.8 ENVIRONMENTAL CONDITIONS 4-7
4.9 SHOCK AND VIBRATION 4-7
4.10 RELIABILITY 4-8
4.11 DISK ERRORS 4-8

Chapter 5
BASIC PRINCIPLES OF OPERATION 5-1
5.1 QUANTUM BIGFOOT CY DRIVE MECHANISM 5-1
5.1.1 Base Casting Assembly 5-3
5.1.2 DC Motor Assembly 5-3
5.1.3 Disk Stack Assemblies 5-3
5.1.4 Headstack Assembly 5-4
5.1.5 Rotary Positioner Assembly 5-4
5.1.6 Automatic Actuator Lock 5-4
5.1.7 Air Filtration 5-4
5.2 DRIVE ELECTRONICS 5-6
5.2.1 µController 5-7
5.2.2 DCIIA 5-7
5.2.3 Read/Write ASIC 5-10
5.2.4 PreAmplifier and Write Driver 5-11
5.3 SERVO SYSTEM 5-12
5.3.1 General Description 5-12
5.3.2 Servo Burst and Track Information 5-12
5.4 READ AND WRITE OPERATIONS 5-13
5.4.1 The Read Channel 5-13
5.4.2 The Write Channel 5-13
5.4.3 Interface Control 5-14
5.5 FIRMWARE FEATURES 5-14
5.5.1 Disk Caching 5-14
5.5.2 Track and Cylinder Skewing 5-16
5.5.3 Error Detection and Correction 5-17
5.5.4 Defect Management 5-23

Chapter 6
IDE BUS INTERFACE AND ATA COMMANDS 6-1

6.1 INTRODUCTION 6-1
6.2 SOFTWARE INTERFACE 6-1
6.3 MECHANICAL DESCRIPTION 6-1
6.3.1 Drive Cable and Connector 6-1
6.4 ELECTRICAL INTERFACE 6-1
6.4.1 IDE Bus Interface 6-1
6.4.2 Host Interface Timing 6-8
6.5 REGISTER ADDRESS DECODING 6-11
6.6 REGISTER DESCRIPTIONS 6-13
6.6.1 Control Block Registers 6-13
6.6.2 Command Block Registers 6-14
6.7 COMMAND DESCRIPTIONS 6-19
6.7.1 Recalibrate–1xh 6-19
6.7.2 Read Sectors–20h (with retry), 21h (without retry) 6-19
6.7.3 Write Sector–30h (with retry), 31h (without retry) 6-20
6.7.4 Read Verify Sectors–40h (with retry), 41h (without retry) 6-20
6.7.5 Format Track–50h 6-20
6.7.6 Seek–7xh 6-21
6.7.7 Execute Drive Diagnostic–90h 6-21
6.7.8 Initialize Drive Parameters–91h 6-22
6.7.9 Read Multiple–C4h 6-22
6.7.10 Write Multiple–C5h 6-22
6.7.11 Set Multiple Mode–C6h 6-23
6.7.12 Read Buffer–E4h 6-23
6.7.13 Write Buffer–E8h 6-23
6.7.14 Power Management Commands 6-24
6.7.15 Identify Drive–ECh 6-26
6.7.16 Set Features–EFh 6-29
6.7.17 Read Defect List 6-30
6.7.18 Configuration 6-32
6.8 ERROR REPORTING 6-37

GLOSSARY G-1

INDEX I-1

List of Figures

FIGURE	DESCRIPTION	PAGE
3-1	Bigfoot CY 2.1 AT Mechanical Dimensions	3-1
3-2	Packaging for a 1-Pack Shipping Container	3-2
3-3	Connector and Jumper Locations	3-3
3-4	Jumper Locations on the Drive Interface Connector	3-4
3-5	Bigfoot CY 2.1/4.3/6.4AT Mounting Dimensions (in millimeters)	3-7
3-6	Mounting Screw Clearance	3-8
3-7	DC Power and IDE-Bus Combination Connector	3-9
3-8	Drive Power Supply and IDE-Bus Interface Cables	3-12
3-9	Drive Installation	3-12
5-1	Quantum Bigfoot CY 4.3 AT Hard Disk Drive Exploded View	5-2
5-2	HDA Air Filtration	5-5
5-3	Quantum Bigfoot CY 2.1/4.3/6.4AT Hard Disk Drive Block Diagram	5-6
5-4	DCIIA Block Diagram	5-7
5-5	Read/Write ASIC Block Diagram	5-10
5-6	Sector Data Field with ECC Check Bytes	5-18
5-7	Byte Interleaving	5-18
5-8	Correctable and Uncorrectable Double-Burst Errors	5-20
5-9	Correctable and Uncorrectable Triple-Burst Errors	5-21
5-10	Nine Correctable Random Burst Errors	5-22
6-1	PIO Interface Timing	6-9
6-2	Multiword DMA Bus Interface Timing	6-10
6-3	Host Interface RESET Timing	6-10

List of Tables

TABLE	DESCRIPTION	PAGE
3-1	Bigfoot CY 2.1/4.3/6.4AT Jumper Options	3-5
3-2	Mating Power Connectors	3-10
3-3	Logical Addressing Format	3-14
4-1	Bigfoot CY 2.1/4.3/6.4AT Specifications	4-1
4-1	Bigfoot CY 2.1/4.3/6.4 AT Specifications (Continued)	4-2
4-2	Formatted Capacity	4-2
4-3	Timing Specifications	4-3
4-4	Power On Reset Thresholds	4-4
4-5	Typical Power and Current Consumption	4-5
4-6	Acoustical Characteristics - Sound Pressure	4-6
4-7	Acoustical Characteristics–Sound Power	4-6
4-8	Mechanical Dimensions	4-6
4-9	Environmental Specifications	4-7
4-10	Shock and Vibration Specifications	4-7
4-11	Error Rates	4-8
5-1	Bigfoot CY 2.1/4.3/6.4 Cylinder Contents	5-3
5-2	Track and Cylinder Skews	5-17
6-1	Drive Connector Pin Assignments (J1, Section C)	6-3
6-2	Relationship of Drive Signals to the IDE Bus	6-7
6-3	PIO Host Interface Timing	6-8
6-4	Multiword DMA Host Interface Timing	6-9
6-5	Host Interface RESET Timing	6-10
6-6	I/O Port Functions and Selection Addresses	6-12
6-7	Command Block Register Initial Values	6-12
6-8	Device Control Register Bits	6-13
6-9	Drive Address Register Bits	6-14
6-10	Error Register Bits	6-15
6-11	Drive Head Register Bits	6-16
6-12	Status Register Bits	6-17
6-13	Quantum Bigfoot CY 2.1/4.3/6.4 AT Command Codes and Parameters	6-18
6-14	Diagnostic Codes	6-21
6-15	Valid Count Range	6-25
6-16	Identify Drive Parameters	6-26
6-17	READ DEFECT LIST LENGTH Command Bytes	6-30
6-18	READ DEFECT LIST Command Bytes	6-31
6-19	DEFECT LIST Data Format	6-31
6-20	DEFECT ENTRY Data Format	6-32
6-21	Accessing the READ CONFIGURATION Command	6-32
6-22	Accessing the SET CONFIGURATION Command	6-33
6-23	Accessing the SET CONFIGURATION WITHOUT SAVING TO DISK Command	6-34
6-24	Configuration Command Format	6-34
6-25	Command Errors	6-37

Chapter 1
ABOUT THIS MANUAL

This chapter gives an overview of the contents of this manual, including the intended audience, manual organization, and terminology and conventions. In addition, it provides a list of other references that might be helpful to the reader.

1.1 AUDIENCE DEFINITION

The Quantum Bigfoot CY 2.1/4.3/6.4AT product manual is intended for several audiences including the original equipment manufacturer (OEM), distributor, installer, and end user. The manual provides you with information about installation, principles of operation, interface command implementation, and maintenance.

1.2 MANUAL ORGANIZATION

This manual provides information about installation, principles of operation, and interface command implementation. It is organized into the following chapters:

- Chapter 1—About This Manual
- Chapter 2—General Description
- Chapter 3—Installation
- Chapter 4—Specifications
- Chapter 5—Basic Principles of Operation
- Chapter 6—IDE-Bus Interface and ATA Commands

In addition, this manual contains a glossary of terms and an index to help you locate important information.

1.3 TERMINOLOGY AND CONVENTIONS

In the Glossary at the back of this manual, you can find definitions for many of the terms used in this manual. In addition, the following abbreviations are used in this manual:

- ASIC application-specific integrated circuit
- bpi bits per inch
- dB decibels
- dBA decibels, A weighted
- ECC error correcting code
- fci flux changes per inch

- Hz hertz
- KB kilobytes
- LSB least significant bit
- mA milliamperes
- MB megabytes (1 MB = 1,000,000 bytes when referring to disk storage and 1,048,576 bytes in all other cases)
- Mbit/s megabits per second
- MB/s megabytes per second
- MHz megahertz
- ms milliseconds
- MSB most significant bit
- mv millivolts
- ns nanoseconds
- tpi tracks per inch
- µs microseconds
- V volts

The typographical and naming conventions used in this manual are listed below. Conventions that are unique to a specific table appear in the notes that follow that table.

Typographical Conventions:

- **Names of Bits:** Bit names are presented in initial capitals. An example is the Host Software Reset Bit.
- **Commands:** Firmware commands are listed as all capitals. An example is WRITE LONG.
- **Parameters:** Parameters are given as initial capitals when spelled out, and are given as all capitals when abbreviated. Examples are Prefetch Enable (PE) and Cache Enable (CE).
- **Hexadecimal Notation:** The hexadecimal notation is given in 9-point subscript form. An example is 30_H.
- **Signal Negation:** A signal name that is defined as active low is listed with a minus sign following the signal. An example is RD–.
- **Messages:** A message that is sent from the drive to the host is listed in all capitals. An example is BUS DEVICE RESET.

Naming Conventions:

- **Host:** In general, the system in which the drive resides is referred to as the host. The AT/IDE host adapter is considered to be part of the host
- **Computer Voice:** This refers to items you type at the computer keyboard. These items are listed in 10-point, all capitals, Courier font. An example is An example is `FORMAT C:/S`.

1.4 REFERENCES

For additional information about the AT interface, refer to:

- IBM Technical Reference Manual #6183355, March 1986
- ATA Common Access Method Specification, Revision 3, January 17, 1995
- Small Form Factor Specification for ATA Timing Extensions #8011, July 2, 1993

Chapter 2
GENERAL DESCRIPTION

This chapter summarizes general functions and key features of the Quantum Bigfoot CY ™ 2.1/4.3/6.4AT hard disk drives, as well as the applicable standards and regulations it meets.

2.1 PRODUCT OVERVIEW

Quantum's Bigfoot CY 2.1/4.3/6.4AT hard disk drives are part of a family of high-performance, 19 mm-high (0.75 in) and 25.4 mm-high (1.0 in), hard disk drives manufactured to meet the highest product quality standards. Bigfoot CY 2.1/4.3/6.4AT hard disk drives use nonremovable, 130 mm (5 1/4-in) hard disks. These drives can be used in IBM®PC/AT-compatible host computer systems, which provide an AT Attachment (ATA) Interface either on an adapter board or on the system motherboard.

Note: Sometimes, the ATA Interface is referred to as the *AT Bus* or *IDE Interface*.

Quantum Bigfoot CY 2.1/4.3/6.4AT hard disk drives feature an embedded IDE drive controller and use ATA commands. The drive manages media defects and error recovery internally, so these operations are transparent to the user.

The innovative design of Quantum Bigfoot CY 2.1/4.3/6.4AT hard disk drives enables Quantum to produce a family of low-cost, high-reliability drives.

2.2 KEY FEATURES

Quantum Bigfoot CY 2.1/4.3/6.4AT hard disk drives include the following key features:

General

- Formatted storage capacity of 2168 MB (1-disk), 4335MB (2-disk), or 6510 MB (3-disk)
- Industry-standard, 130 mm (5 1/4-in) form factor
- Low-profile, 19 mm (0.75 in), height (1-disk) or 25.4 mm (1.00 in) height (2- and 3-disk)
- Emulation of IBM®PC/AT®task file register and all ATA fixed-disk commands
- Embedded servo design
- Fast ATA-2 support
- 16/17 PRML recording code
- MR heads

Performance

- Data transfer rates of up to 16 MB/s using PIO Mode 4 with IORDY or 16 MB/s using multiword DMA Mode 2.
- Average seek time of 12 ms (1-disk) and 14 ms (2- and 3-disk)
- Average rotational latency of 8.33ms
- 1:1 interleave on read/write operations
- 128K buffer with 70 Kbyte adaptive segmentation cache. Look-ahead DisCache ® feature with continuous prefetch, and WriteCache™write-buffering capabilities
- Support for all ATA data transfer modes with PIO Mode 4 and multiword DMA mode 2
- Auto Write, Auto Read, and Auto Transfer

Reliability

- Power-Up Self diagnostic firmware
- 224-bit, interleaved Reed-Solomon Error Correcting Code (ECC), with cross-checking and double- and triple-burst correction for bursts up to 96 bits in length
- Double-burst ECC correction on-the-fly
- 300,000 hour Mean Time Between Failure (MTBF) in the field
- Automatic retry on read errors
- 20,000 Contact Start/Stop cycles (controlled)
- Transparent media-defect mapping
- Defect-free interface
- Reassignment of defective sectors discovered in the field, without reformatting
- High-performance, in-line defective sector skipping
- Patented AIRLOCK®automatic shipping lock and dedicated landing zone
- Failure prediction
- S.M.A.R.T. 4 system

Versatility

- Power-saving modes
- Downloadable firmware (DiskWare)
- Cable select feature
- Common CD-ROM cable ready
- Ability to daisy-chain two drives on the interface
- Multi-Sector Auto Read and Multiple Auto Write transfer
- LBA mode addressing/Extended CHS mode

2.3 STANDARDS AND REGULATIONS (Pending)

Quantum Bigfoot CY 2.1/4.3/6.4AT hard disk drives satisfy the following standards and regulations under 5.25 Series Safety name:

- Tested to CISPR22 EN55022-4, which meets Federal Communications Commission (FCC) Rules for Radiated and Conducted Emissions, Part 15, Sub Part J, for Class B Equipment in an enclosure.
- Underwriters Laboratory (U.L.): Standard 1950. Information technology equipment including business equipment.
- Canadian Standards Association (CSA): Standard C22.2 No. 950. Information technology equipment including business equipment.
- European Standards: Technischer Überwachungs Verein (TÜV). Standard EN 60 950 and IEC 950. Information technology equipment including business equipment in an enclosure.
- EMC testing to meet EU Directives (89/336/EEC) for Electromagnetic Compatibility and outlined in EN50082-1 and EN50082-2.

Note: The Standards and Regulations information is pending.

Chapter 3
INSTALLATION

This chapter explains how to unpack, configure, mount, and connect the Bigfoot CY 2.1/4.3/6.4AT hard disk drive prior to operation. It also explains how to start up and operate the drive.

3.1 SPACE REQUIREMENTS

Quantum ships the Bigfoot CY 2.1/4.3/6.4AT hard disk drives without a faceplate. Figure 3-1 shows the external dimensions of the Bigfoot CY 2.1 AT drive. The Bigfoot CY 2.1AT drive has a thickness of 19 mm (0.75 in.). See Figure 3-6 and Figure 3-7 for a view of the Bigfoot CY 4.3/6.4AT drives, which are thicker (1-inch). The other dimensions of the Bigfoot CY 4.3/6.4AT hard disk drives are the same as those of the Bigfoot CY 2.1 AT drive, shown below.

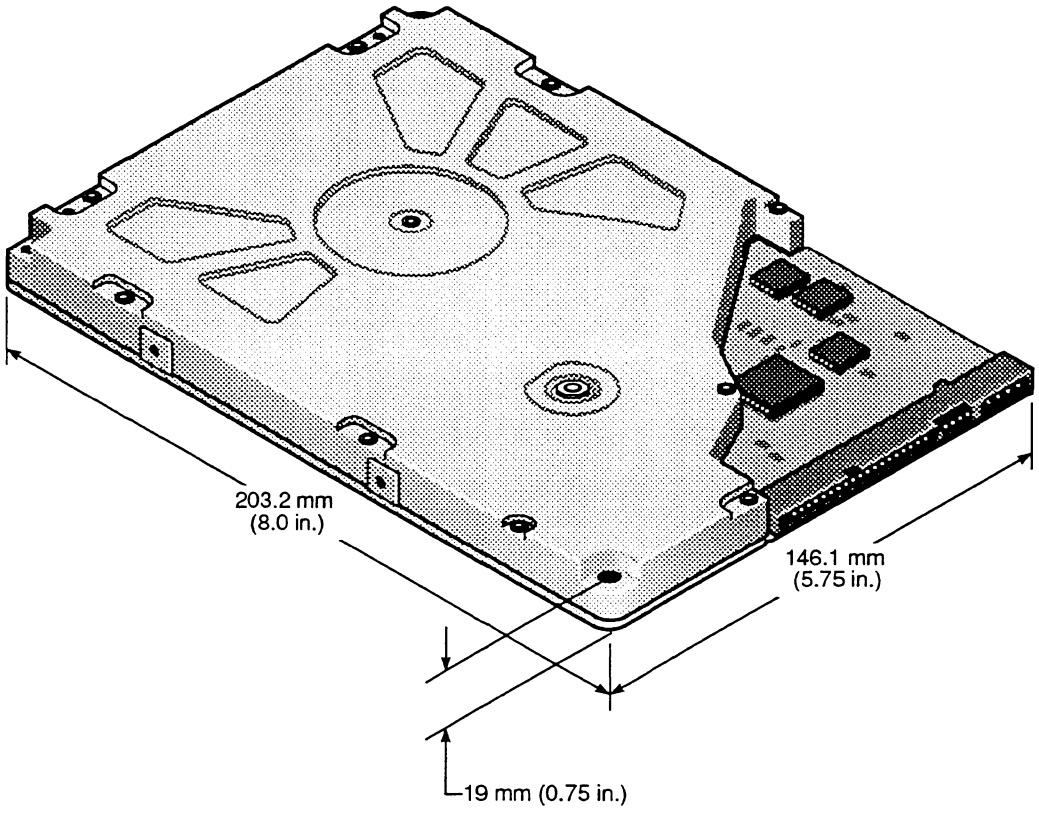

Figure 3-1 *Bigfoot CY 2.1 AT Mechanical Dimensions*

3.2 UNPACKING INSTRUCTIONS

CAUTION: The maximum limits for physical shock can be exceeded if the drive is not handled properly. Special care should be taken not to bump or drop the drive.

1. Open the shipping container.
2. Remove the upper protective packaging pad from the box.

   Figure 3-2 shows the packaging for the Bigfoot CY 2.1/4.3/6.4AT hard disk drives in a 1-pack shipping container. (A 10-pack shipping container is available for multiple drive shipments and has the same type of protective packaging.)

3. Remove the drive from the box.

CAUTION: During shipment and handling, the drive is packed in an antistatic electrostatic discharge (ESD) bag to prevent electronic component damage to ESD sensitive devices. Remove the drive from the ESD bag only when you are ready to install it. To avoid accidental damage to the drive, do not use a sharp instrument to open the ESD bag.

4. Save the packaging materials for possible future use.

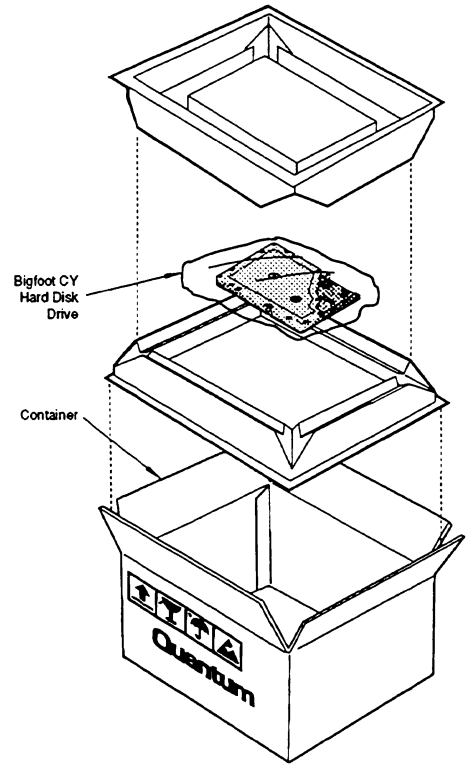

Figure 3-2 *Packaging for a 1-Pack Shipping Container*

3.3 JUMPER SETTINGS

The configuration of a Bigfoot CY 2.1/4.3/6.4AT hard disk drives depends on the host system into which it is to be installed. This section describes the jumper setting options you must take into account prior to installation. Figure 3-3 shows the connector and jumper locations.

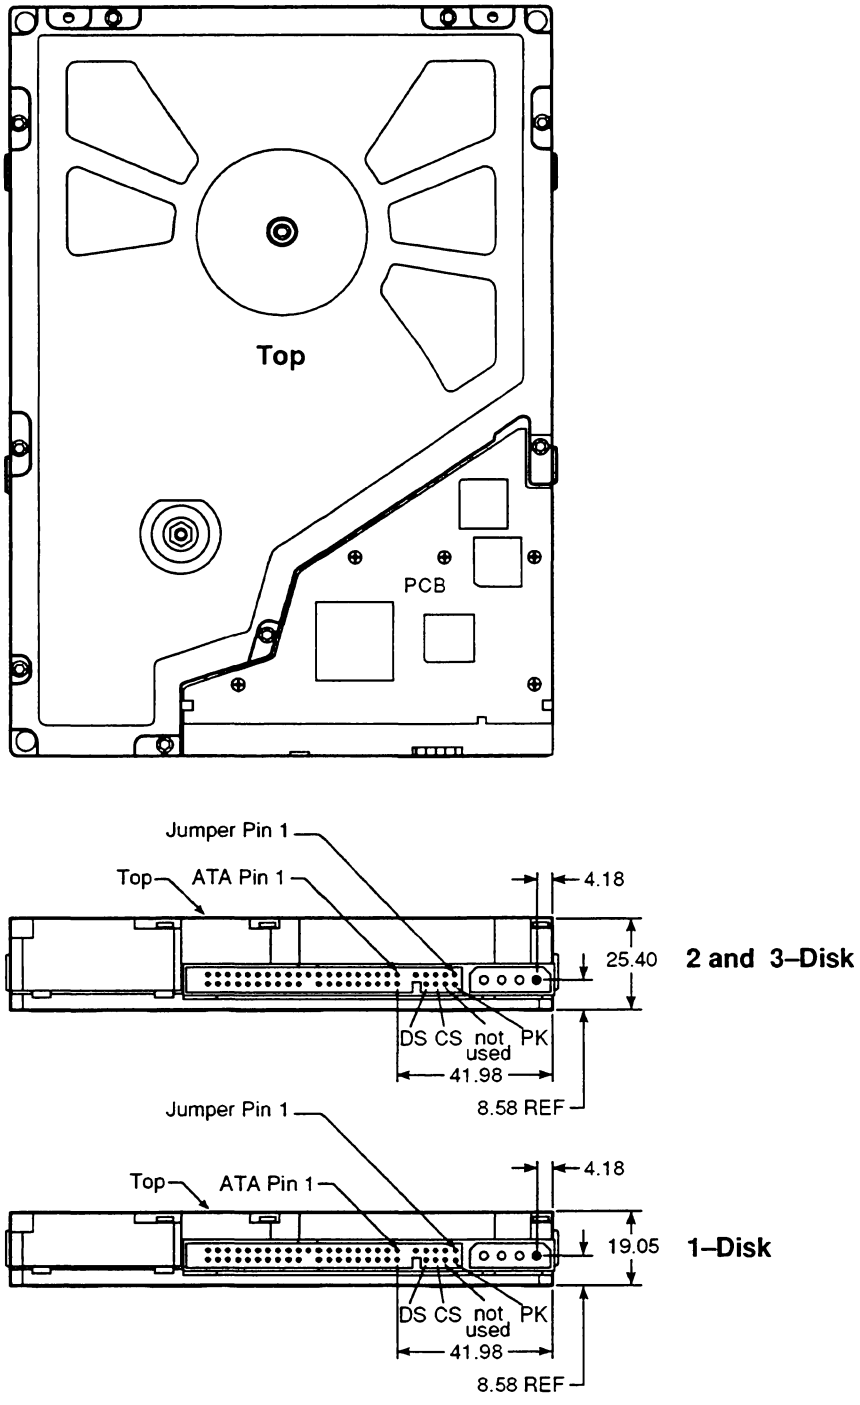

Figure 3-3 *Connector and Jumper Locations*

Figure 3-4 shows the details of the interface connector, which has an embedded 8-pin jumper connector.

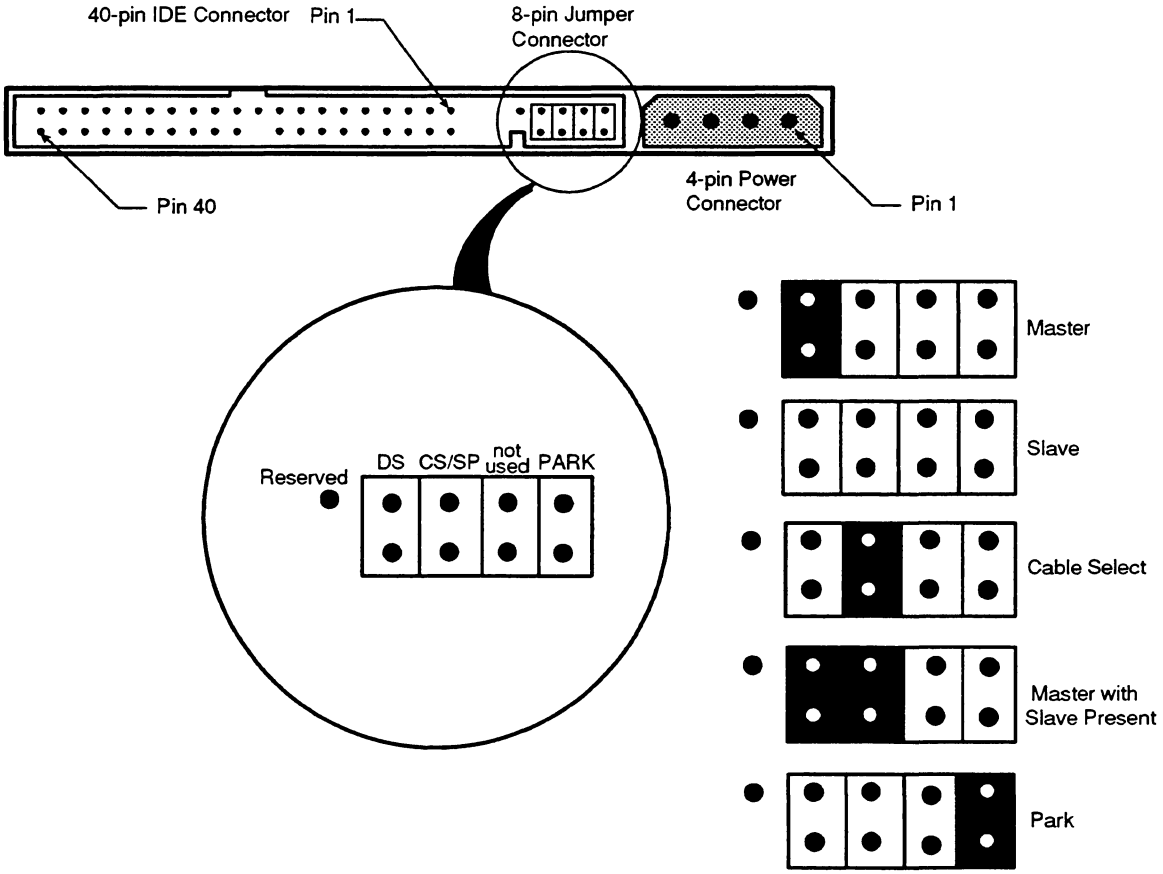

Figure 3-4 *Jumper Locations on the Drive Interface Connector*

The 8-pin embedded connector has two jumper locations, DS and CS/SP, that provide a way to configure the drive's mode of operation:

- DS–Drive Select
- CS/SP–Cable Select/Slave Present

A "PARK" position is also provided to store a jumper when it is not in use.

Table 3-1 defines the jumper configurations for the drive and the operational modes that result from the configurations.

Table 3-1 *Bigfoot CY 2.1/4.3/6.4AT Jumper Options*

CS/SP	DS	PARK	Cable Pin 28	DESCRIPTION
0	0	X	X	Drive is configured as a Slave
0	1	X	X	Drive is configured as a Master
1	1	X	X	Drive is configured as a Master, with an attached Slave
1	X	X	open	Drive is configured as a Slave
1	X	X	ground	Drive is configured as a Master

Note: In Table 3-1, a "0" indicates that the jumper is removed, a "1" indicates that the jumper is installed, and X indicates the jumper setting does not matter.

3.3.1 Cable Select/Slave Present (CS/SP) Jumper

When the CS jumper is installed, the drive uses pin 28 of the interface connector to determine if the drive is a master or a slave. If pin 28 of the interface connector is grounded, the drive is configured as a Master; if pin 28 of the connector is left open, an internal pullup brings the pin high, and the drive is configured as a Slave. If the CS and DS jumpers are both installed, the drive is configured as a Master with an attached Slave.

3.3.2 Drive Select (DS) Jumper

For systems that do not support the Cable Select feature, a drive can be configured as a Master or Slave by removing the CS jumper and using only the DS jumper. If the DS jumper is installed, the drive is configured as a Master; if it is removed, the drive is configured as a Slave.

3.3.3 Park Jumper

The PARK position is used as a place to store a jumper for a Slave drive in systems that do not support Cable Select.

3.4 IDE-BUS ADAPTER

There are two ways you can configure a system to allow the Bigfoot CY 2.1/4.3/6.4AT hard disk drives to communicate over the IDE-bus of an IBM or IBM-compatible PC:

1. Connect the drive to a 40-pin IDE-bus connector (if available) on the motherboard of the PC.

2. Install an IDE-compatible adapter board in the PC and connect the drive to the adapter board.

3.4.1 40-Pin IDE-Bus Connector

Many of the later design PC motherboards have a built-in 40-pin, IDE-bus connector, that is compatible with the 40-pin IDE interface of the Bigfoot CY 2.1/4.3/6.4AT hard disk drives. If the motherboard has an IDE connector, simply connect a 40-pin ribbon cable between the drive and the motherboard.

You should also refer to the motherboard instruction manual, and refer to Chapter 6 of this manual to ensure signal compatibility.

3.4.2 Adapter Board

If your PC motherboard does not contain a built-in, 40-pin IDE-bus interface connector, you must install an IDE-bus adapter board and connecting cable to allow the drive to interface with the motherboard. Quantum does not supply such an adapter board, but they are available from several third-party vendors.

Please read over carefully the instruction manual for the adapter board you purchase, as well as Chapter 6 of this manual to ensure signal compatibility between the adapter board and the drive. Also, make sure that the adapter board jumper settings are appropriate.

3.5 MOUNTING

Drive mounting orientation, clearance, and ventilation requirements are described in the following subsections. For mounting, M3 screws are required. To avoid stripping the mounting-hole threads, the maximum torque applied to the screws must not exceed 5 kg-cm.

3.5.1 Orientation

The mounting holes on the Bigfoot CY 2.1/4.3/6.4AT hard disk drives allow the drive to be mounted in any orientation.

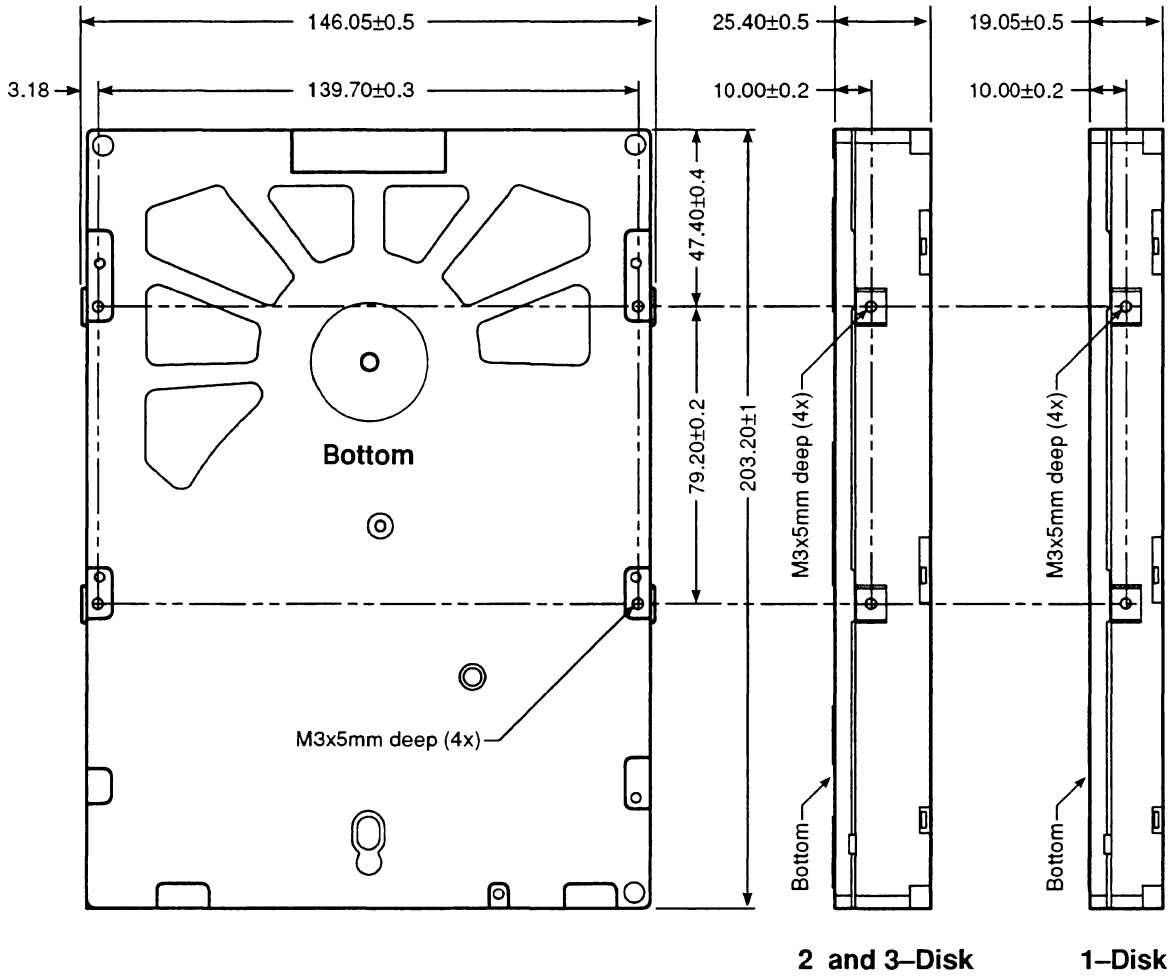

Figure 3-5 *Bigfoot CY 2.1/4.3/6.4AT Mounting Dimensions (in millimeters)*

3.5.2 Mounting Screw Clearance

The printed-circuit board assembly (PCBA) is very close to the mounting holes. Figure 3-6 shows the 1-inch version of the Bigfoot CY 2.1/4.3/6.4AT and specifies the clearance between the screws in the mounting holes and the PCBA. Do not use mounting screws longer than the maximum lengths specified in Figure 3-6. The specified screw length allows full use of the mounting-hole threads, while avoiding damaging or placing unwanted stress on the PCBA.

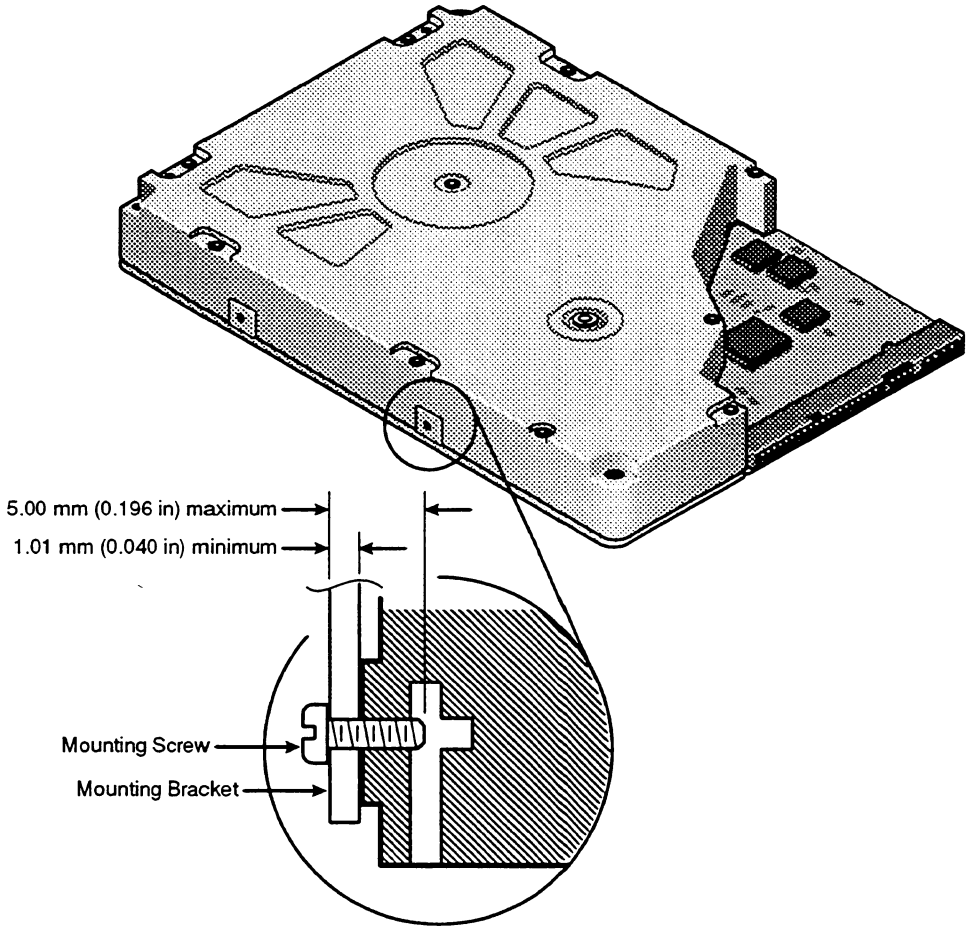

Figure 3-6 *Mounting Screw Clearance*

CAUTION: The PCBA is very close to the mounting holes. Do not exceed the specified length for the mounting screws. The specified screw length allows full use of the mounting-hole threads, while avoiding damaging or placing unwanted stress on the PCBA. Figure 3-6 specifies the minimum clearance between the PCBA and the screws in the mounting holes. To avoid stripping the mounting-hole threads, the maximum torque applied to the screws must not exceed 5 kg-cm.

3.5.3 Ventilation

The Bigfoot CY 2.1/4.3/6.4AT hard disk drives operates without a cooling fan, provided the ambient air temperature does not exceed 55°C (131°F).

3.6 COMBINATION CONNECTOR

As shown in Figure 3-7, there is a three-section combination connector mounted on the back edge of the PCBA. The connector contains a 40-pin IDE interface, an 8-pin jumper block, and a 4-pin power connector.

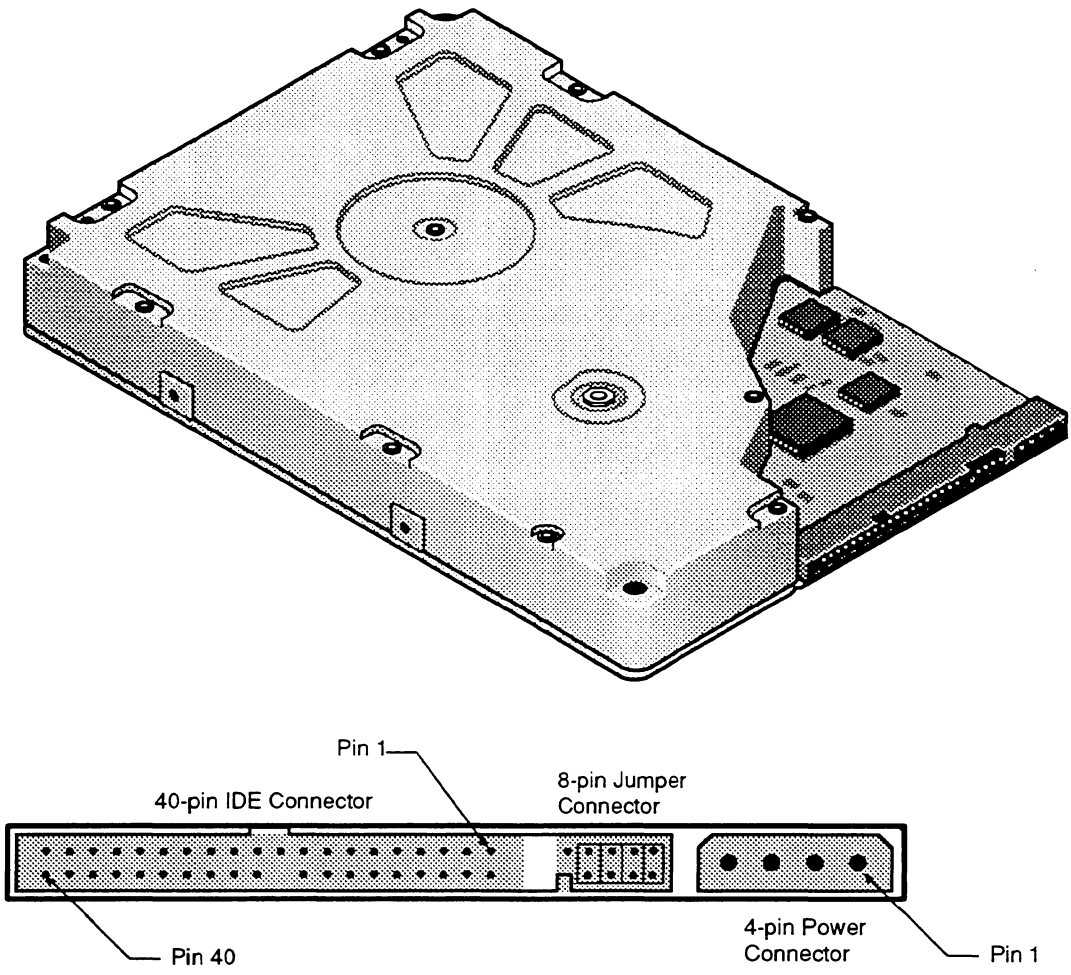

Figure 3-7 *DC Power and IDE-Bus Combination Connector*

3.6.1 DC Power

The recommended mating connectors for the +5 Vdc and +12 Vdc input power are listed in Table 3-2.

Table 3-2 *Mating Power Connectors*

	PIN NUMBER	VOLTAGE LEVEL	MATING CONNECTOR TYPE AND PART NUMBER (OR EQUIVALENT)
4-Pin Connector	1	+12 Vdc	4-Pin Connector: AMP P/N 1-480424-0
	2	+12 Vdc Return (Ground)	Loose piece contacts: AMP P/N 61173-4
	3	+5 Vdc Return (Ground)	Strip contacts: AMP P/N 350078-4
	4	+5 Vdc	

3.6.2 IDE-Bus Interface Connector

On the Bigfoot CY 2.1/4.3/6.4AT hard disk drives, the IDE-bus interface cable connector is a 40-pin Universal Header, as shown in Figure 3-7.

See Chapter 6, "IDE-Bus Interface and ATA Commands," for more detailed information about the required signals. Refer to Table 6-1 for the pin assignments of the IDE-bus connector.

For mating with the 40-pin connector, recommended cable connectors include the following parts or their equivalents:

AMP receptacle with strain relief	P/N 1-499506-0
AMP receptacle without strain relief	P/N 1-746193-0

To key the 40-pin cable connector, you must plug the hole that corresponds to pin 20.

Other recommended part numbers for the mating connector include:

40-Pin Connector	3M 3417-7000 or equivalent
Strain Relief	3M 3448-2040 or equivalent
Flat Cable (Stranded 28 AWG)	3M 3365-40 or equivalent
Flat Cable (Stranded 28 AWG)	3M 3517-40 (shielded) or equivalent

3.7 DRIVE INSTALLATION

You can install the Bigfoot CY 2.1/4.3/6.4AT hard disk drives in an AT-compatible system in two ways:

3.7.1 For Systems With A Motherboard IDE Adapter

You can install the Bigfoot CY 2.1/4.3/6.4AT hard disk drives in an AT-compatible system that contains a 40-pin IDE-bus connector on the motherboard.

To connect the drive to the motherboard, use a 40-pin ribbon cable. Ensure that pin 1 of the drive is connected to pin 1 of the motherboard connector.

3.7.2 For Systems With An IDE Adapter Board

To install a Bigfoot CY 2.1/4.3/6.4AT hard disk drives in an AT-compatible system without a 40-pin IDE-bus connector on its motherboard, you need a third-party IDE-compatible adapter board.

3.7.2.1 Adapter Board Installation

Carefully read the manual that accompanies your adapter board before installing it. Make sure that all the jumpers are set properly and that there are no addressing or signal conflicts. Install the adapter board in your system according to the adapter board manual.

3.7.2.2 Connecting the Adapter Board and the Drive

Use a 40-pin ribbon cable to connect the drive to the board. See Figure 3-8. To connect the drive to the board:

1. Insert the 40-pin cable connector into the mating connector of the adapter board. Make sure that pin 1 of the connector matches with pin 1 on the cable.

2. Insert the other end of the cable into the header on the drive. When inserting this end of the cable, make sure that pin 1 of the cable connects to pin 1 of the drive connector.

3. Secure the drive to the system chassis by using the mounting screws, as shown in Figure 3-9.

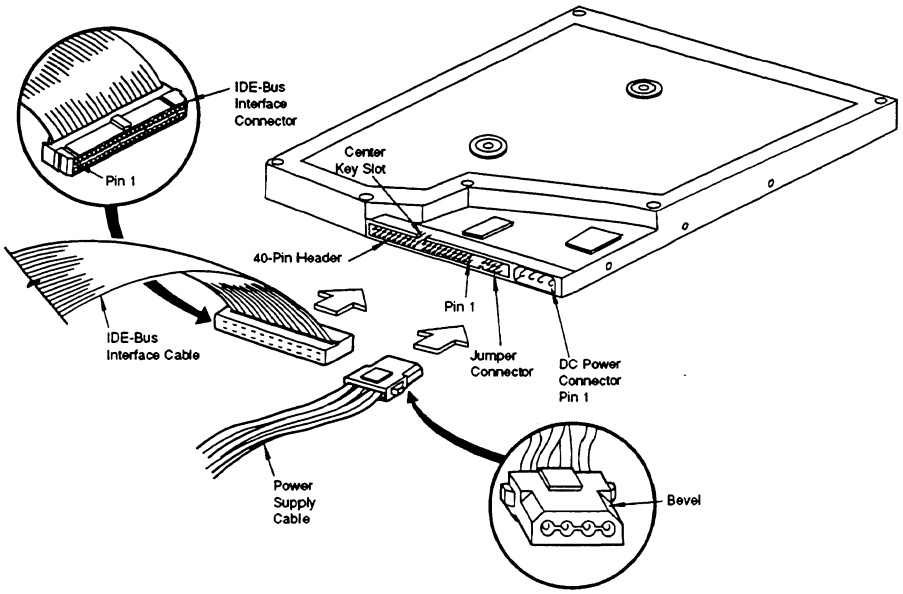

Figure 3-8 *Drive Power Supply and IDE-Bus Interface Cables*

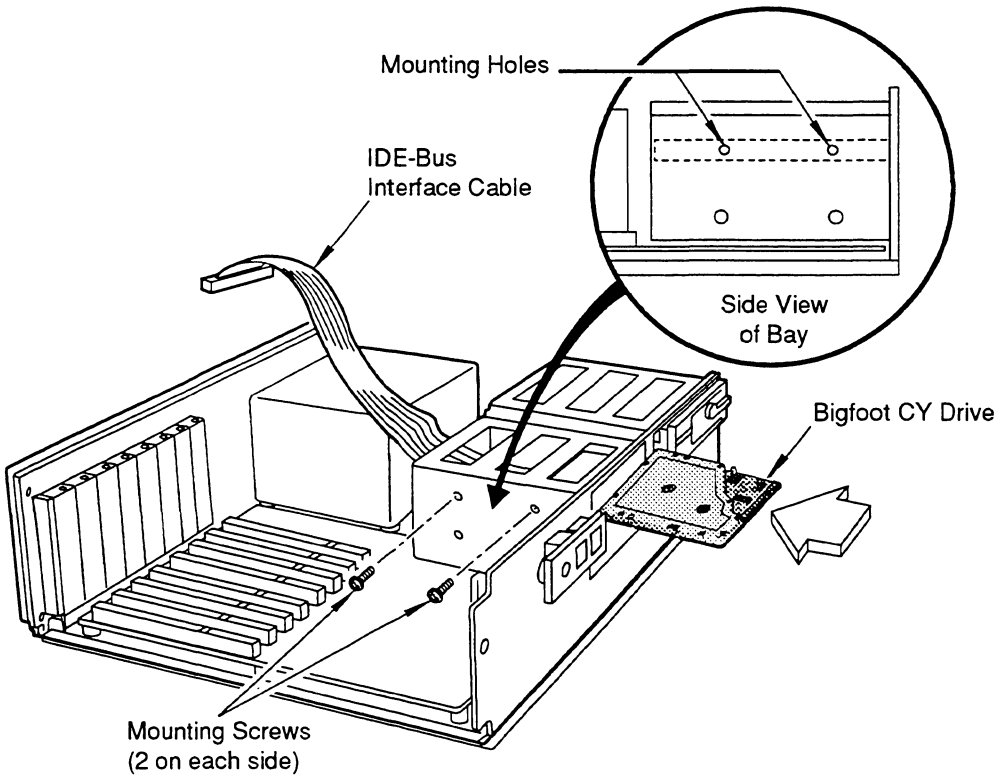

Figure 3-9 *Drive Installation*

3.8 TECHNIQUES IN DRIVE CONFIGURATION

3.8.1 Maximum Partition Size in MS-DOS 4.0 and Above

When using MS-DOS 4.0 and above, FDISK allows the user to partition a hard drive with a total capacity of 4 GB (4,294,967,296 bytes). However, the file allocation table (FAT) only supports partition sizes of 2 GB (2,147,483,648 bytes). Therefore, a hard drive between the sizes of 2 and 4 GB must be divided into multiple partitions, each not exceeding 2 GB.

3.8.2 1024 Cylinder Limitation on Older Computer Systems

Because MS-DOS uses the computer's ROM BIOS to access the hard drive, the cylinder limitation of 1024 from the AT ROM BIOS will only be viewed by the MS-DOS Operating System. The CMOS System Setup is able to scan the total number of cylinders, but the BIOS is still limited to use only 1024 cylinders. The following are some techniques to resolve this difficulty.

- Use a third party software program that translates the hard drive parameters to an acceptable configuration for MS-DOS.
- Use a hard disk controller that translates the hard drive parameters to an appropriate setup for both MS-DOS and the computer system's ROM BIOS.

3.8.3 Newer Computer Systems with Extended BIOS Translation

Some newer computer systems allow the user to configure disk drives that go beyond the 528 MB (528,482,304 bytes) barrier. Here are formulas to translate drives with a maximum capacity of 8.4 GB (8,422,686,720 bytes):

xcyl = cyl * n	xcyl is defined as a new cylinder translation
xhead = head * n	xhead is defined as a new head translation
xsec = sec = 63	xsec is defined as a new sector translation

where n = 2, 4, 8, ..., a power of 2

n is chosen to reduce the number of cylinders to be less than or equal to 1024. However, sectors must equal 63 and the number of heads cannot exceed 255.

Note: Be advised that the previous information is dependent upon the capabilities of the computer system, hard disk controller, and/or software programs. Some configurations may not provide the user proper operation of the disk drive. All other documentation should be examined prior to the hard drive installation.

3.9 SYSTEM STARTUP AND OPERATION

Once you have installed the Bigfoot CY 2.1/4.3/6.4AT hard disk drive and adapter board (if required) in the host system, you are ready to partition and format the drive for operation. To set up the drive correctly, follow these steps:

1. Power on the system.

2. Run the SETUP program. This is generally on a Diagnostics or Utilities disk, or within the system's BIOS.

The SETUP program allows you to enter the types of optional hardware installed–such as the hard disk drive type, the floppy disk drive capacity, and the display adapter type. The system's BIOS uses this information to initialize the system when the power is switched on. For instructions on how to use the SETUP program, refer to the system manual for your PC.

3. Enter the appropriate parameters.

During the AT system CMOS setup, you must enter the drive type for the Bigfoot CY 2.1/4.3/6.4AT hard disk drives. This procedure allows the system to recognize the drive by translating its *physical* drive geometry parameters such as cylinders, heads, and sectors per track, into a *logical* addressing mode.

Bigfoot CY 2.1/4.3/6.4AT drives can work with different BIOS drive-type tables of various host systems. You can choose any drive type that does not exceed the capacity of the drive. Table 3-3 gives the logical parameters that provide the maximum capacity on Bigfoot CY 2.1/4.3/6.4AT hard disk drives.

Table 3-3 *Logical Addressing Format*

	Bigfoot CY 2.1AT	**Bigfoot CY 4.3AT**	**Bigfoot CY 6.4AT**
Logical Cylinders	4,092	8,960	13,446
Logical Heads	16	15	15
Logical Sectors/Track	63	63	63
Total Number Logical Sectors	4,124,736	8,467,200	12,706,470
Logical Capacity	2,111,864,832	4,335,206,400	6,505,712,640

To match the logical specifications of the drive to the drive type of a particular BIOS, consult the system's drive-type table. This table specifies the number of cylinders, heads, and sectors for a particular drive type.

4. Boot the system using the operating system installation disk–for example, *MS-DOS*–then follow the installation instructions in the operating system manual.

Note: You do not need to perform a low-level format on the drive because it was done at the factory before shipment.

Chapter 4
SPECIFICATIONS

This chapter gives a detailed description of the physical, electrical, and environmental characteristics of the Bigfoot CY 2.1/4.3/6.4AT hard disk drives.

4.1 SPECIFICATION SUMMARY

Table 4-1 gives a summary of the Bigfoot CY 2.1/4.3/6.4AT hard disk drives.

Table 4-1 *Bigfoot CY 2.1/4.3/6.4AT Specifications*

DESCRIPTION	Bigfoot CY 2.1AT	Bigfoot CY 4.3AT	Bigfoot CY 6.4AT
Capacity (formatted)	2,111.86 MB	4,335.20 MB	6505.71 MB
Nominal rotational speed (rpm)	3,600	3,600	3,600
Number of disks	1	2	3
Number of R/W heads	2	4	6
Data Organization			
Zones per surface	15	15	15
Tracks per surface	8,895	8,895	8,895
Total tracks	17,790	35,580	53,370
Sectors per track			
Inside zone	170	170	170
Outside zone	288	288	288
Total User Sectors	4,124,736	8,467,200	12,706,470
Bytes per sector	512	512	512
Number of tracks per cylinder	2	4	6
Recording			
Recording technology	Multiple Zone	Multiple Zone	Multiple Zone
Maximum linear density	135,806 fci	135,806 fci	135,806 fci
Maximum Recording Density	127,817 (bpi)	127,817 (bpi)	127,817 (bpi)
Encoding method	PRML 16/17	PRML 16/17	PRML 16/17
Interleave	1:1	1:1	1:1
Track density (avg.)	6,950 tpi	6,950 tpi	6,950 tpi
Effective areal density	888 Mbit/in^2	888 Mbit/in^2	888 Mbit/in^2

Table 4-1 *Bigfoot CY 2.1/4.3/6.4 AT Specifications (Continued)*

DESCRIPTION		Bigfoot 2.1AT	Bigfoot 4.3/6.4AT
Performance			
Typical Seek times[1]			
Average seek		12.0 ms	14.0 ms
Track-to-track sequential		3.0 ms	3.0 ms
Full stroke		25 ms	25 ms
Sequential Head Switch		3.0 ms	3.0 ms
Rotational Latency		8.33 ms	8.33 ms
Data Transfer rate			
Disk to Read Buffer[2]		55.63 to 92.61 Mb/s[3]	55.63 to 92.61 Mb/s[4]
Read/Write Buffer to IDE-bus (PIO Mode)	Without IORDY	6.0 MB/s max.	6.0 MB/s max.
	With IORDY	16.67 MB/s max.	16.67 MB/s max.
Read/Write Buffer to IDE-bus (DMA Mode)		16.67 MB/s max.	16.67 MB/s max.
Buffer size (The upper 32K is used for firmware)		128 KB	128 KB
Projected MTBF		300,000 hrs	300,000 hrs
Contact Start/Stop Cycles[5]		20,000	20,000
Auto head-park method		AirLock	AirLock

1. Seek times are at nominal conditions and include settling for read.
2. Disk to read buffer transfer rate is zone-dependent.
3. Megabits per second.
4. Megabits per second.
5. CSS specifications assume a duty cycle of one power off operation for every four idle spin-down.

4.2 FORMATTED CAPACITY

At the factory, the Bigfoot CY 2.1/4.3/6.4AT hard disk drives receive a low-level format that creates the actual tracks and sectors on the drive. Table 4-2 shows the capacity resulting from this process. For operation with DOS, UNIX, or other operating systems, formatting done at the user level results in less capacity than the physical capacity shown in Table 4-2.

Table 4-2 *Formatted Capacity*

	Bigfoot CY 2.1AT	Bigfoot CY 4.3AT	Bigfoot CY 6.4AT
Number of 512-byte sectors	4,124,736	8,467,200	12,706,470

4.3 DATA TRANSFER RATES

Data is transferred from the disk to the read buffer at a rate up to 6.5 MB/s in bursts. Data is transferred from the read buffer to the AT bus at a rate of up to 6.0 MB/s, using programmed I/O without IORDY. If IORDY is used, then this transfer rate can be increased to 16.67 MB/s. Refer to paragraph 6.6.15 IDENTIFY DRIVE command for additional information. Using the Multiword DMA protocol, transfer rates of 16.67 MB/s. are achievable (burst mode only).

4.4 TIMING SPECIFICATIONS

Table 4-3 illustrates the timing specifications of the Bigfoot CY 2.1/4.3/6.4AT hard disk drive.

Table 4-3 *Timing Specifications*

Parameter	Bigfoot CY 2.1/4.3/6.4AT	
	Typical Nominal[1]	Maximum Worst Case[2]
Single Track Seek[3]	3.0 ms	3.5 ms
Sequential Cylinder Switch Time[4]	3.0 ms	3.5 ms
Sequential Head Switch Time[5]	3.0 ms	3.5 ms
Random Average (Read or Seek)[6]	12.0 ms/14.0 ms/14.0 ms	12.5 ms/14.5 ms/14.5 ms
Random Average (Write)	17.0 ms	18 ms
Full-Stroke Seek	25 ms	25 ms
Average Rotational Latency	8.33 ms	–
Power On[7] to Drive Ready[8]	8/11/13.5 seconds	10/14/18 seconds
Power On to Interface Ready[9]	11.5 seconds	18.0 seconds
Standby[10] to Drive Ready	4.5/7.5/9.5 seconds	7/11/15 seconds
Drive Ready to Shutdown	15/20/25 seconds	20/26/32 seconds[11]

1. Nominal conditions are as follows:
   - Nominal temperature (25° C)
   - Nominal supply voltages (12.0V, 5.0V)
   - No applied shock or vibration
2. Worst-case conditions are as follows:
   - Worst-case temperature extremes (4°C to 71°C)
   - Worst-case supply voltages (12V ±10%, 5V±5%)
3. Seek time is defined as the time required for the actuator to seek and settle on-track. It is measured by averaging 5000 seeks of the indicated type as shown in this table. The seek times listed include head settling time, but do not include command overhead time or rotational latency delays.
4. Sequential Cylinder Switch Time is the time from the conclusion of the last sector of a cylinder to the first logical sector on the next cylinder.
5. Sequential Head Switch Time is the time from the last sector of a track to the beginning of the first logical sector of the next track of the same cylinder.
6. Average seek time is the average of 5000 random seeks. When a seek error occurs, recovery for that seek can take up to seven seconds.
7. Power On is the time from when the supply voltages reach operating range to when the drive is ready to accept any command.
8. Drive Ready is the condition in which the disks are rotating at the rated speed and the drive is able to accept and execute commands requiring disk access without further delay.
9. Interface Ready is the condition in which the drive is ready to accept any command, before the disks are rotating at rated speed.
10. Standby is the condition in which the microprocessor is powered, but not the HDA. When the host sends the drive a standby command, the drive parks the heads away from the data zone and spins down to a complete stop.
11. After 20/26/32 seconds, it is safe to move the disk drive.

4.5 POWER

The Bigfoot CY 2.1/4.3/6.4Bigfoot CY 2.1/4.3/6.4AT hard disk drives operate from two supply voltages:

- +12V ±10%
- +5V ±5%

The allowable ripple and noise for each voltage is 250 mV for the +12 Vdc supply and 150 mV for the +5 Vdc supply.

4.5.1 Power Sequencing

You can apply the power in any order or manner, or short or open either the power or power- return line with no loss of data or damage to the disk drive. However, data may be lost in the sector being written at the time of power loss. The drive can withstand transient voltages of +10% to -10% from nominal while powering up or down.

4.5.2 Power Reset Limits

When powering up, the drive remains reset (inactive) until both V_{HT} reset limits are exceeded. When powering down, the drive becomes reset when either supply voltage drops below the V_{LT} threshold. See Table 4-4.

Table 4-4 *Power On Reset Thresholds*

DC VOLTAGE	THRESHOLD		THRESHOLD	
	+5V		+12V	
	Max	Min	Max	Min
V_{LT}[1]	4.65 V[2]	4.15 V	9.35 V[1]	8.30 V
V_{HT}[1]	4.75 V	4.25 V[1]	9.50 V	8.45 V[1]
H_{YST}[1]	130 mV (typ)	130 mV (typ)	250 mV (typ)	250 mV (typ)

1. V_{LT} and V_{HT} values are rounded to the nearest 0.05 V; H_{YST} values are rounded to the nearest 5 mV.
2. Includes a 100 mV Peak-Peak ripple on 5V or 250 mV Peak-Peak ripple on 12V to maximize or minimize values.

4.5.3 Power Requirements

Table 4-5 lists the voltages and corresponding current for the various modes of operation of the Bigfoot CY 2.1/4.3/6.4AT hard disk drives.

Table 4-5 *Typical Power and Current Consumption*

MODE OF OPERATION	Typical Average Current (mA rms)[8]						Typical Average Power (Watts)		
	Bigfoot CY 2.1AT		Bigfoot CY 4.3AT		Bigfoot CY 6.4AT				
	+12V	+5V	+12V	+5V	+12V	+5V	2.1AT	4.3AT	6.4AT
Startup[1] (peak)	2000	600	2000	600	2000	600	27.0	27.0	27.0
Maximum Seek[2]	750	550	900	550	900	550	11.8	13.6	13.6
Operating[3]	400	550	600	550	675	550	7.6	10.0	10.9
Read/Write on Track[4]	300	600	375	600	425	600	6.6	7.5	8.1
Idle[5]	300	500	375	500	425	500	6.1	7.0	6.2
Standby[6]	10	325	10	325	10	325	1.4	1.4	1.4
Sleep[7]	10	250	10	250	10	250	1.0	1.0	1.0

Notes:

1. Startup is stated as the peak (>10 ms) power required during spindle startup. This power will be required for less than 6 seconds.
2. Maximum Seek is for continuous random seek operations with 8.3 ms delay between seek complete indication and the next seek command.
3. Operating mode is defined when data is being read from or written to the disk. It is computed based on 40% seeking, 30% on-track reading and 30% on-track writing.
4. Read/Write/Ontrack is defined as 50% read operations and 50% write operations on a single physical track.
5. Idle is when the drive is not reading, writing, or seeking, the motor is up to speed and the DRIVE READY condition exists. Actuator is residing on last track accessed.
6. Standby is when the motor is stopped, actuator parked, and all electronics except the interface control is in a low power state. STANDBY will occur after a programmable time-out from the last host access. Drive ready and seek complete status exist. The drive will leave STANDBY upon receipt of a command that requires disk access, or upon receipt of a spin-up command.
7. Sleep is when the spindle and actuator motors are off with the heads latched in the landing zone. The microprocessor has also entered the halt or power-down state. Receipt of a reset causes the drive to transition from the sleep to the standby mode.
8. Current is rms. (except for startup)

4.6 ACOUSTICS

Table 4-6 and Table 4-7 specifies the acoustical characteristics of the Bigfoot CY 2.1/4.3/6.4AT hard disk drives.

Table 4-6 *Acoustical Characteristics - Sound Pressure*

OPERATING MODE	MEASURED NOISE	DISTANCE
Idle On Track	34 dbA (typical) 39 dbA (maximum)	1 m (39.3 in) 1 m (39.3 in)
Seeking Random	38 dbA (typical) 42 dbA (maximum)	1 m (39.3 in) 1 m (39.3 in)

Table 4-7 *Acoustical Characteristics–Sound Power*

Operating Mode	Measured Noise (Sound Power per ISO 7779)
Idle On Track	3.9 Bels (typical) 4.2 Bels (maximum)
Seeking Random	4.5 Bels (typical) 4.8 Bels (maximum)

4.7 MECHANICAL DIMENSIONS

Table 4-8 specifies the mechanical dimensions of the Bigfoot CY 2.1/4.3/6.4AT hard disk drives. Dimensions measurements do not include the faceplate.

Table 4-8 *Mechanical Dimensions*

Dimension	Millimeters	Inches
Height		
Bigfoot CY 2.1	19.0 mm	0.75 in
Bigfoot CY 4.3	25.4 mm	1.00 in
Bigfoot CY 6.4	25.4 mm	1.00 in
Width	146.2 mm	5.75 in
Depth	203.2 mm	8.0 in

Mechanical Weight

Model	Grams	Pounds
Bigfoot CY 2.1	909	2.0
Bigfoot CY 4.3	1,091	2.4
Bigfoot CY 6.4	1,136	2.5

4.8 ENVIRONMENTAL CONDITIONS

Table 4-9 summarizes the environmental specifications for the Bigfoot CY 2.1/4.3/6.4AT hard disk drive.

Table 4-9 *Environmental Specifications*

PARAMETER	OPERATING	NONOPERATING
Temperature[1]	0°C to 55°C 32° C to 131°F	-40°C to 75°C -40°C to 167°F
Temperature Gradient[1]	24°C/hr 75.2 °F/hr	48°C/hr 118.4 °F/hr
Humidity[1] Maximum Wet Bulb	5% to 85% rh 29°C (84.2°F)	5% to 95% rh 46°C (115°F)
Humidity Gradient	30% hr	30% hr
Altitude[2]	-200 m to 3,000 m (-650 to 10,000 ft.)	-200 m to 12,000 m (-650 to 40,000 ft.)
Altitude Gradient	1.5 kPa/min.	8 kPa/min.

1. No condensation
2. Altitude is relative to sea level

4.9 SHOCK AND VIBRATION

The Bigfoot CY 2.1/4.3/6.4AT hard disk drives can withstand levels of shock and vibration applied to any of its three mutually perpendicular axes, or principal base axes, as specified in Table 4-10. A functioning drive can be subjected to specified *operating* levels of shock and vibration. When a drive has been subjected to specified *nonoperating* levels of shock and vibration, with power to the drive off, there will be no change in performance at power on.

When packed in either the 1-pack or 10-pack shipping container, Bigfoot CY 2.1/4.3/6.4AT drives can withstand a drop from 30 inches onto a concrete surface on any of its surfaces, six edges, or three corners.

Table 4-10 *Shock and Vibration Specifications*

	OPERATING	NONOPERATING
Shock	10G, 1/2 sine wave, 11 ms 18 G, 1/2 sine wave, 3 ms	70G, 11ms, 1/2 sine wave 110 G, 3 ms 1/2 sine wave 200 G, 2 ms, 1/2 sine wave
Vibration[1] 1.0 octave/minute sweep	1G p-p, 5-300 Hz	2G p-p, 5-300 Hz

1. No unrecovered errors.

4.10 RELIABILITY

Mean Time Between Failures (MTBF):	The projected field MTBF is 300,000 Power On Hours (POH). The Quantum MTBF numbers represent Bell-Core TR-TSY-000332 MTBF predictions and represent the minimum MTBF that Quantum or a customer would expect from the drive.
Preventive Maintenance (PM):	Not required
Start/Stop:	20,000 cycles (minimum)
Component Life:	5 Years

4.11 DISK ERRORS

Table 4-11 provides the error rates for the Bigfoot CY 2.1/4.3/6.4AT hard disk drives.

Table 4-11 *Error Rates*

ERROR TYPE	MAXIMUM NUMBER OF ERRORS
Reallocated data errors:	
Correctable read errors[1]	1 error per 10^{12} bits read
Uncorrectable read errors[2]	1 error per 10^{14} bits read
Seek Errors[3]	1 error per 10^{6} seeks

Notes:

1. Correctable read errors are read errors that are recovered by retries or by application of the triple-burst error correction algorithm.
2. Uncorrectable read errors are errors that are not correctable using ECC or retries. The drive will terminate retry reads either when a repeating error pattern occurs, or after eight unsuccessful retries and the application of triple-burst error correction.
3. Seek errors occur when the actuator fails to reach (or remain) over the requested cylinder or if the drive executes a recalibration routine to find the requested cylinder (a full recalibration takes about seven seconds).

Chapter 5
BASIC PRINCIPLES OF OPERATION

This chapter describes the operation of Quantum Bigfoot CY 2.1/4.3/6.4AT functional subsystems. It is intended as a guide to the operation of the drive, rather than a detailed theory of operation.

5.1 QUANTUM BIGFOOT CY DRIVE MECHANISM

This section describes the drive mechanism. Section 5.2 describes the drive electronics. The Quantum Bigfoot CY 2.1/4.3/6.4AT hard disk drives consist of a mechanical assembly and a PCB as shown in Figure 5-1. The drawing is illustrated with two disks, showing the Quantum Bigfoot CY 4.3 AT drive configuration. The Quantum Bigfoot CY 6.4 AT hard disk drive contains three hard disks, and the Quantum Bigfoot CY 2.1AT hard disk drive contains one hard disk.

The head/disk assembly (HDA) contains the mechanical subassemblies of the drive, which are sealed under a metal cover. The HDA consists of the following components:

- Base casting
- DC motor assembly
- Disk stack assembly
- Headstack assembly
- Rotary positioner assembly
- Automatic actuator lock
- Air filter

The drive is assembled in a Class-100 clean room.

CAUTION: To ensure that the air in the HDA remains free of contamination, never remove or adjust its cover and seals. Tampering with the HDA will void your warranty.

The Quantum Bigfoot CY 2.1/4.3/6.4AT drives are a one-, two-, or three-disk product family. The Quantum Bigfoot CY 2.1AT hard disk drive contains one magnetic disk and two read/write heads. The Quantum Bigfoot CY 4.3AT drives contain two magnetic disks and four read/write heads, and the Quantum Bigfoot CY 6.4AT drive contains three magnetic disks and six read/write heads

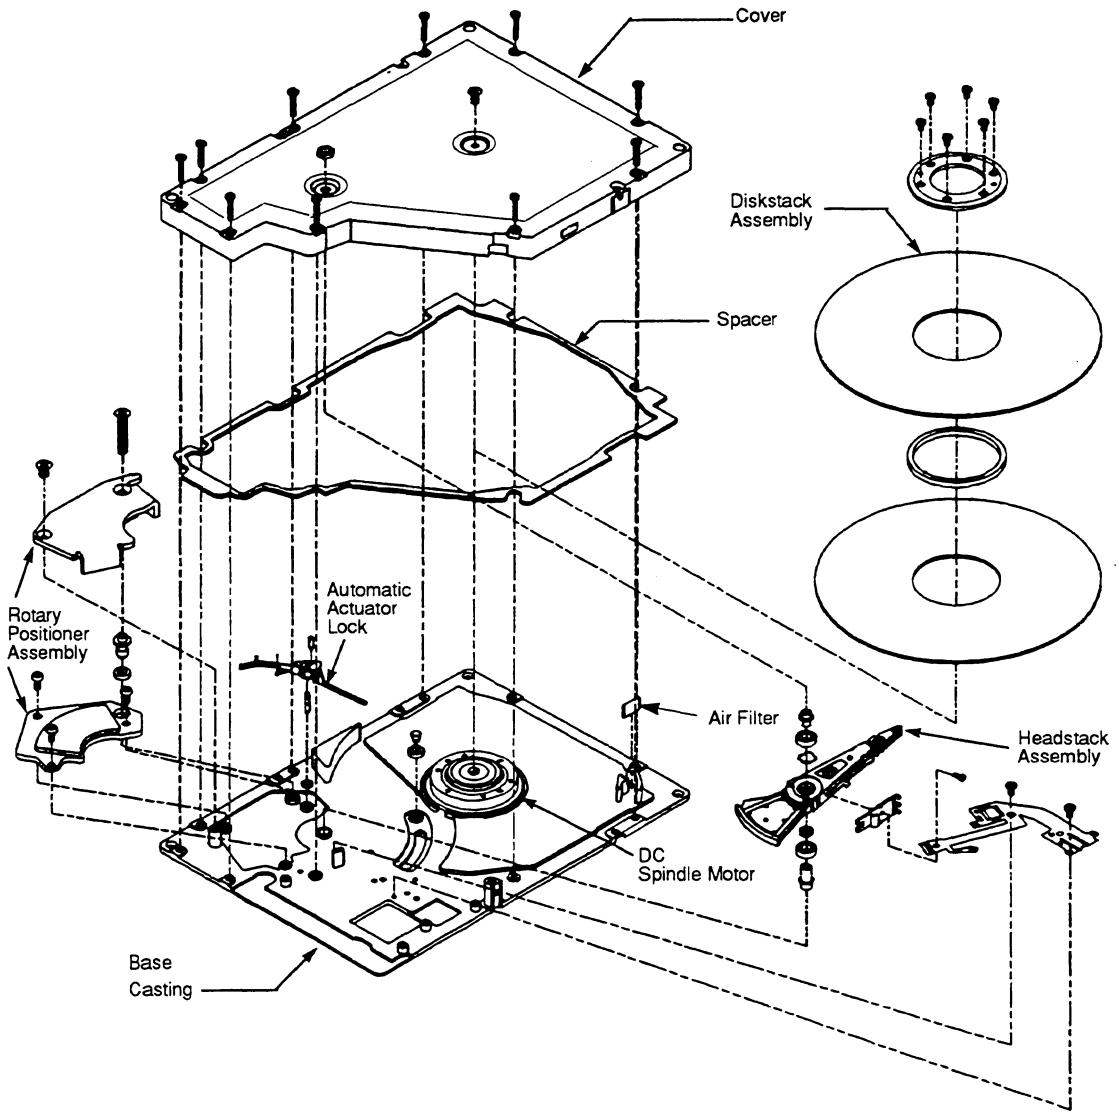

Figure 5-1 *Quantum Bigfoot CY 4.3 AT Hard Disk Drive Exploded View*

5.1.1 Base Casting Assembly

A single-piece, aluminum-alloy base casting provides a mounting surface for the drive mechanism and PCB. The base casting also acts as the flange for the DC motor assembly. To provide a contamination-free environment for the HDA, a gasket provides a seal between the base casting and the metal cover that encloses the drive mechanism.

5.1.2 DC Motor Assembly

Integral with the base casting, the DC motor assembly is a fixed-shaft, brushless DC spindle motor that drives the counter-clockwise rotation of the disks.

5.1.3 Disk Stack Assemblies

The disk stack assembly in the Quantum Bigfoot CY 2.1AT hard disk drive consists of one disk secured by a disk clamp. The Quantum Bigfoot CY 4.3AT hard disk drive contains two disks, and the Quantum Bigfoot CY 6.4AT hard disk drive contains three disks. The disks have an overcoat of sputtered hydrogenated carbon or other Quantum-approved material.

A carbon overcoat lubricates the disk surface, to prevent head and media wear due to head contact with the disk surface during head takeoff and landing. Head contact with the disk surface occurs only in the landing zone outside of the data area, when the disk is not rotating at full speed. The landing zone is located at the inner diameter of the disk, beyond the last cylinder of the data area.

The Quantum Bigfoot CY 2.1/4.3/6.4AT hard disk drives have 8,941 tracks per recording surface. Of these tracks, 46 are used for system data, leaving 8,895 for data. The data tracks are divided into 15 recording zones. The drive uses multiple zone recording, where each data track contains between 156 and 224 sectors, depending on the recording zone. The sectors per track allocation for each zone is provided in Table 5-1.

Table 5-1 *Bigfoot CY 2.1/4.3/6.4 Cylinder Contents*

CYLINDER CONTENT	ZONE[1]	NUMBER OF TRACKS	CYLINDER AREA	SECTORS PER TRACK	DATA RATE (Mbit/s)
System Data	0	20	0	224	70.496
User Data	1	593	0 - 592	304	98.635
	2	593	593 - 1185	302	98.188
	3	593	1186 - 1778	299	97.322
	4	593	1779 - 2371	288	93.365
	5	593	2372 - 2964	275	89.412
	6	593	2965 - 3557	264	85.752
	7	593	3558 - 4150	252	81.412
	8	593	4151 - 4743	241	78.306
	9	593	4744 - 5336	229	74.458
	10	593	5337 - 5929	224	70.496
	11	593	5930 - 6522	205	66.453
	12	593	6523 - 7115	193	62.830
	13	593	7116 - 7708	180	58.643
	14	593	7709 - 8301	168	53.743
	15	593	8302 - 8894	156	50.824

1. For user data, zone 15 is the innermost zone and zone 1 is the outermost zone.

5.1.4 Headstack Assembly

The headstack assembly consists of read/write heads, an E-block and coil joined together by insertion molding to form an E-block/coil subassembly, bearings, and a flex circuit. Read/write heads mounted to spring-steel flexures are swage mounted onto the rotary positioner assembly arms. The E-block is a single piece, die-cast design.

The flex circuit exits the HDA between the base casting and the cover. A cover gasket seals the gap. The flex circuit connects the headstack assembly to the PCB. The flex circuit contains a read preamplifier/write driver IC.

5.1.5 Rotary Positioner Assembly

The rotary positioner, or rotary voice-coil actuator, is a Quantum-proprietary design that consists of an upper permanent magnet plate and lower flux plate bolted to the base casting, a rotary single-phase coil molded around the headstack mounting hub, and a bearing shaft. The single bi-polar magnet consists of two alternating poles and is bonded to the magnet plate. A resilient crash stop prevents the heads from being driven into the spindle or off the disk surface

Current from the power amplifier induces a magnetic field in the voice coil. Fluctuations in the field around the permanent magnet cause the voice coil to move. The movement of the voice coil positions the heads over the requested cylinder.

5.1.6 Automatic Actuator Lock

To ensure data integrity and prevent damage during shipment, the drive uses a dedicated landing zone and Quantum's patented Airlock®. The Airlock holds the headstack in the landing zone whenever the disks are not rotating. It consists of an air vane mounted near the perimeter of the disk stack and a locking arm that restrains the actuator arm assembly.

When DC power is applied to the motor and the disk stack rotates, the rotation generates an airflow on the surface of the disk. As the flow of air across the air vane increases with disk rotation, the locking arm pivots away from the actuator arm, enabling the headstack to move out of the landing zone. When DC power is removed from the motor, an electronic return mechanism automatically pulls the actuator into the landing zone, where the Airlock holds it in place.

5.1.7 Air Filtration

The Quantum Bigfoot CY 2.1/4.3/6.4AT hard disk drives are Winchester-type drives. The heads fly very close to the media surface. Therefore, it is essential that the air circulating within the drive be kept free of particles. Quantum assembles the drive in a Class-100, purified air environment, then seals the drive with a metal cover. When the drive is in use, the rotation of the disks forces the air inside of the drive through an internal filter.

The highest air pressure within the HDA is at the outer perimeter of the disks. A constant stream of air flows through a 0.3-micron circulation filter positioned in the base casting. As illustrated in Figure 5-2, air flows through the circulation filter in the direction of disk rotation. This design provides a continuous flow of filtered air when the disks rotate.

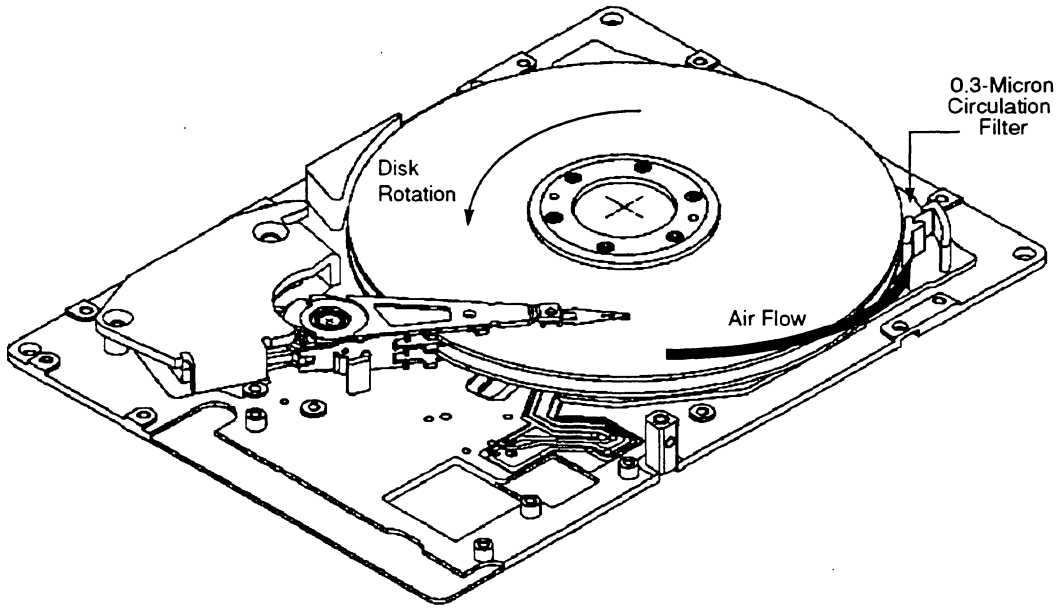

Figure 5-2 *HDA Air Filtration*

5.2 DRIVE ELECTRONICS

Advanced circuit design and the use of miniature surface-mounted devices and proprietary VLSI components enable the drive electronics, including the IDE bus interface, to reside on a single printed circuit board assembly (PCBA).

Note: The components are mounted on only one side of the PCB.

Figure 5-3 contains a simplified block diagram of the Quantum Bigfoot CY 2.1/4.3/6.4AT hard disk drive electronics.

The only electrical component not on the PCBA is the PreAmplifier and Write Driver IC. It is on the flex circuit (inside of the sealed HDA). Mounting the preamplifier as close as possible to the read/write heads improves the signal-to-noise ratio. The flex circuit (including the PreAmplifier and Write Driver IC) provides the electrical connection between the PCB, the rotary positioner assembly, and read/write heads.

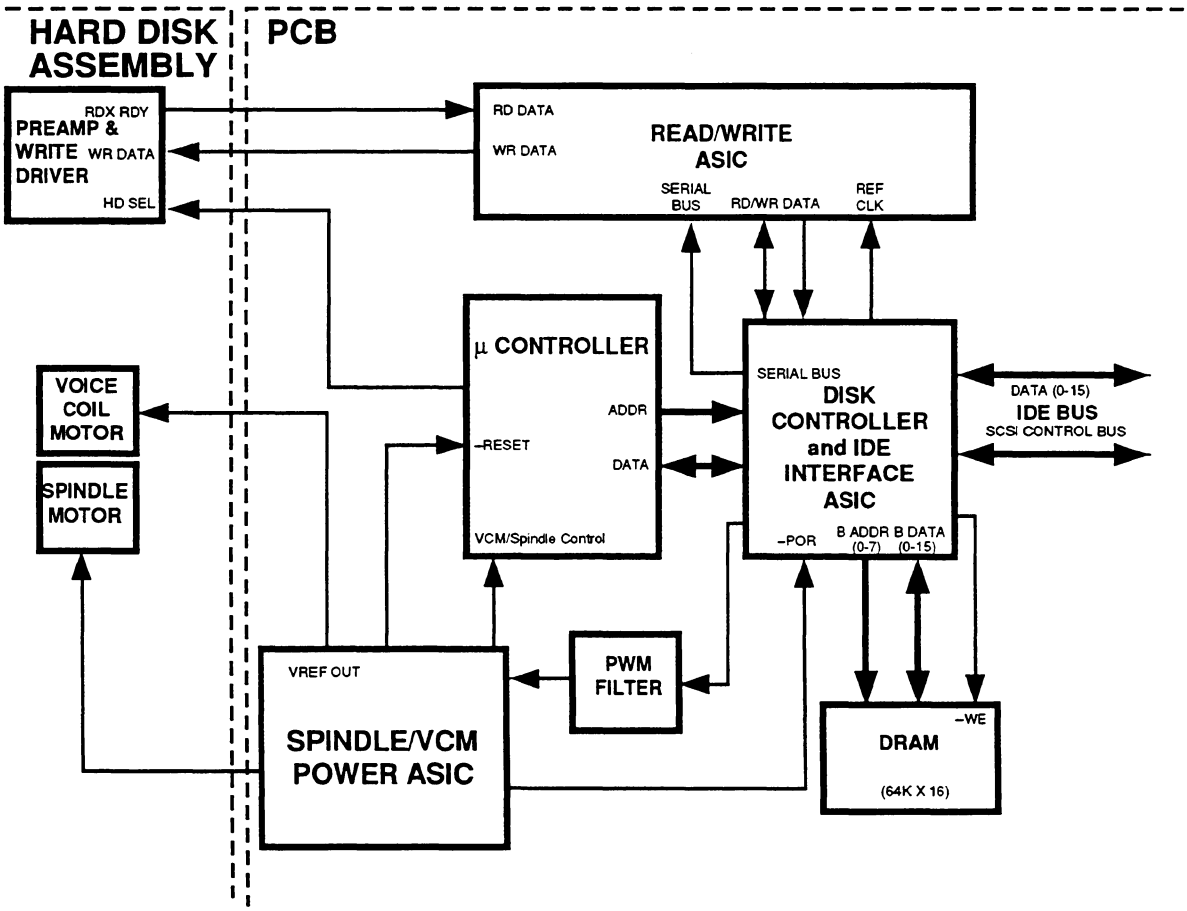

Figure 5-3 *Quantum Bigfoot CY 2.1/4.3/6.4AT Hard Disk Drive Block Diagram*

5.2.1 µController

The µController provides local processor services to the drive electronics under program control. The µController manages the resources of the Disk Controller and IDE Interface ASIC (DCIIA), the Read/Write ASIC, and the Spindle/VCM Driver. In addition, it controls the head selection process.

An internal 54 Kbyte ROM contained within the µController provides program code that the µController executes to complete a drive spinup and recalibration procedure, after which the µController reads additional control code from the disk (diskware) and stores it in the buffer DRAM.

5.2.2 DCIIA

The DCIIA (Disk Controller and IDE Interface ASIC) shown in Figure 5-4 provides control functions to the drive under the direction of the µController.

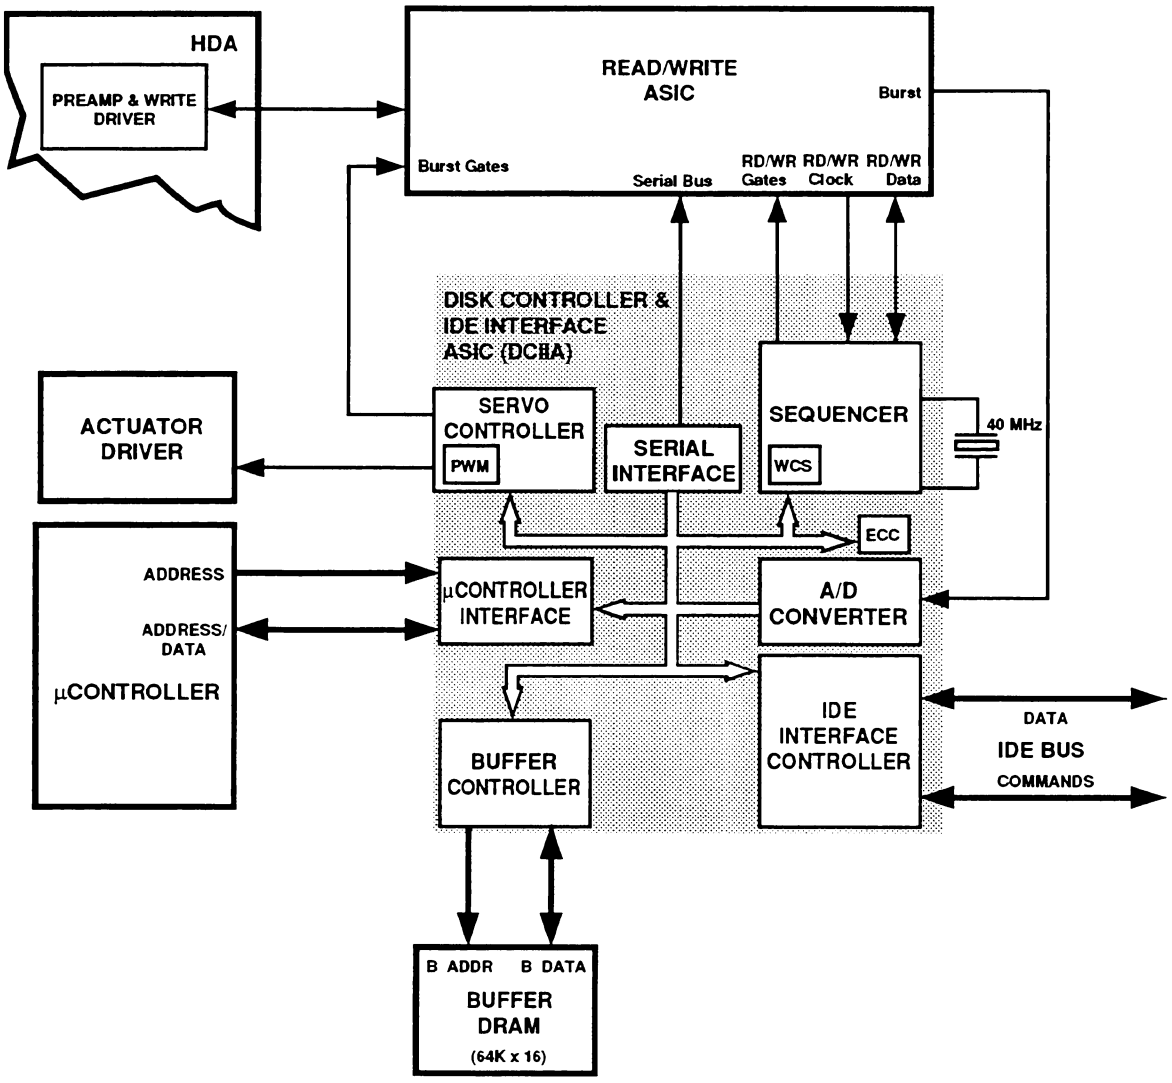

Figure 5-4 *DCIIA Block Diagram*

The DCIIA is a proprietary ASIC developed by Quantum. The DCIIA has eight functional modules (described below):

- 8-bit A/D Converter
- Error Correction Control
- Sequencer
- Buffer Controller
- µController Interface
- Servo Controller, including PWM
- Serial Interface
- IDE Interface Controller

5.2.2.1 A/D Converter

The Analog to Digital converter (A/D) receives multiplexed burst analog inputs from the Read/Write ASIC. The A/D is used to sample the demodulated position information (burst inputs) and convert it to a digital signal the Servo Controller uses to position the HDA actuator.

5.2.2.2 Error Correction Control

The Error Correction Control block utilizes a Reed-Solomon encoder/decoder circuit that is used for disk read/write operations. It uses a total of 20 redundancy bytes organized as 18 ECC (Error Correction Code) bytes and two cross-check bytes. The ECC uses eight bits per symbol and three interleaves. This allows triple-burst error correction of at least 65, and as many as 72 bits in error.

5.2.2.3 Sequencer

The sequencer controls the operation of the read and write channel portions of the DCIIA. To initiate a disk operation, the µController loads a set of commands into the WCS (writable control store) register. Loading and manipulating the WCS is done through the µController Interface registers.

The sequencer also directly drives the read and write gates (RG, WG) of the Read/Write ASIC and the R/W Preamplifier, as well as passing write data to the Precompensator circuit in the Read/Write ASIC.

5.2.2.4 Buffer Controller

The buffer controller supports a 128 Kbyte buffer, which is organized as 64 K x 16 bits. The 16-bit width implementation provides a 32 MB/s maximum buffer bandwidth, which allows 11 MB/s maximum disk channel bandwidth. This increased bandwidth allows the µController to have direct access to the buffer, eliminating the need for a separate µController RAM IC.

The buffer controller supports both drive and host address rollover and reloading, to allow for buffer segmentation. Drive and host addresses may be separately loaded for automated read/write functions.

The Buffer Controller operates under the direction of the µController.

5.2.2.5 µController Interface

The µController Interface provides the means for the µController to read and write data to the DCIIA modules to control their operation, or supply them with needed information. It consists of both physical and logical components.

The physical component of the interface comprises the 16-bit MAD (Multiplexed Address/Data) bus, four additional address lines, read and write strobe, an address latch enable (ALE) signal, and a wait control line.

The logical component of the interface comprises internal control and data registers accessible to the µController. By writing and reading these registers, the µController loads the Sequencer, controls and configures the Buffer controller, and passes coded servo information to the Servo Controller.

5.2.2.6 Servo Controller

The Servo Controller contains a 13-bit Digital to Analog converter (D/A), in the form of a Pulse Width Modulator (PWM). The PWM signal is output to the Actuator Driver to control the motion of the actuator. The Servo Controller also decodes raw data from the disk to extract the current position information. The position information is read by the µController, and is used to generate the actuator control signal that is sent to the PWM. The actuator driver is an analog power amplifier circuit external to the DCIIA. The Servo Controller operates under the direction of the µController.

5.2.2.7 Serial Interface

The Serial Interface provides a high speed Read/Write interface path to the Read/Write ASIC under the direction of the µController.

5.2.2.8 IDE Interface Controller

The IDE Interface Controller portion of the DCIIA provides data handling, bus control, and transfer management services for the IDE interface. Configuration and control of the interface is accomplished by the µController across the MAD bus. Data transfer operations are controlled by the DCIIA Buffer Controller module.

5.2.3 Read/Write ASIC

The Read/Write ASIC shown in Figure 5-5 provides write data precompensation and read channel processing functions for the drive. The Read/Write ASIC receives the RD GATE signal, reference oscillator, serial programming, and servo burst and sample gates from the DCIIA. The Read/Write ASIC sends decoded read data and the read reference clock, and receives write data from the DCIIA. This a highly integrated circuit which is completely under digital control from the DCIIA.

The Read/Write ASIC comprises 11 main functional modules (described below):

- Pre-Compensator
- Variable Gain Amplifier (VGA)
- Butterworth Filter
- FIR Filter
- Flash A/D Converter
- Viterbi Detector
- ENDEC
- Servo Detector and Sample/Hold
- Clock Synthesizer
- PLL
- Serial Interface

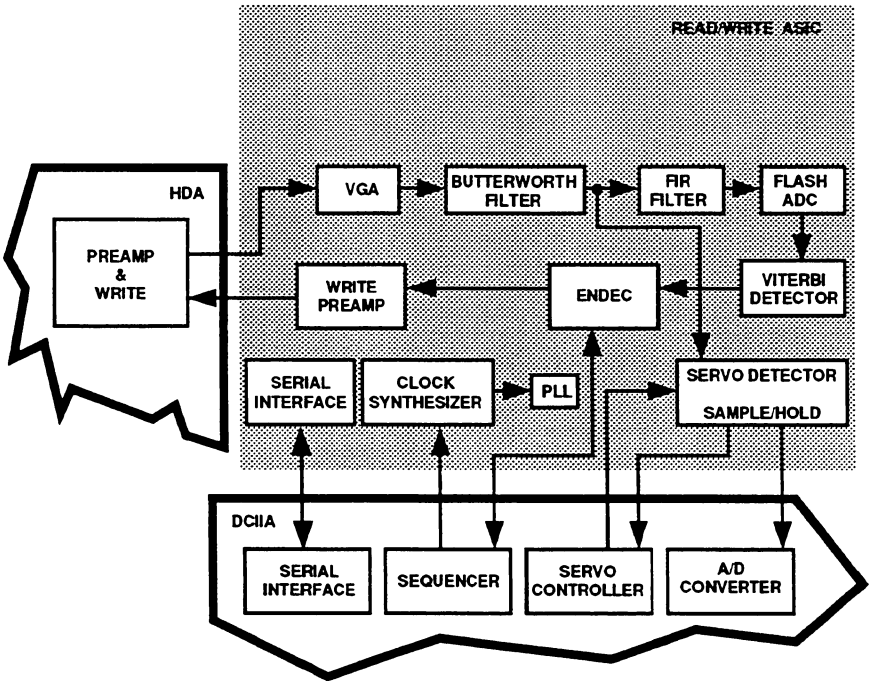

Figure 5-5 *Read/Write ASIC Block Diagram*

5.2.3.1 Pre-Compensator

The pre-compensator introduces pre-compensation to the write data received from the sequencer module in the DCIIA. The pre-compensated data is then passed to the R/W Pre-Amplifier and written to the disk. Pre-compensation reduces the write interference from adjacent write bit.

5.2.3.2 Variable Gain Amplifier (VGA)

Digital and analog controlled AGC function with input attenuator for extended range.

5.2.3.3 Butterworth Filter

Continuous time data filter which can be programmed for each zone rate.

5.2.3.4 FIR Filter

Digitally controlled and programmable filter for partial response signal conditioning.

5.2.3.5 Flash A/D Converter

Provides very high speed digitization of the processed read signal.

5.2.3.6 Viterbi Detector

Decodes ADC result into binary bit stream.

5.2.3.7 ENDEC

Provides 16/17 code conversion to NRZ. Includes preamble and sync mark generation and detection.

5.2.3.8 Servo Detector and Sample/Hold

Peak detection with weighted averaging and multiple sample and hold of servo bursts.

5.2.3.9 Clock Synthesizer

Provides programmable frequencies for each zone data rate.

5.2.3.10 PLL

Provides digital read clock recovery.

5.2.3.11 Serial Interface

High speed interface for digital control of all internal blocks.

5.2.4 PreAmplifier and Write Driver

The PreAmplifier And Write Driver provides write driver and read pre-amplifier functions, and R/W head selection. The write driver receives precompensated write data from the PreCompensator module in the Read/Write ASIC. The write driver then sends this data to the heads in the form of a corresponding alternating current. The read pre-amplifier amplifies the low-amplitude differential voltages generated by the R/W heads and transmits them to the VGA module in the Read/Write ASIC. Head select is determined by the µController.

5.3 SERVO SYSTEM

5.3.1 General Description

The servo system controls the positioning of the read/write heads and holds the read/write heads on track during read/write operations. The servo system also compensates for thermal offsets between heads on different surfaces, and any shock and vibration the drive is subjected to.

The Quantum Bigfoot CY 2.1/4.3/6.4AT hard disk drives use a sectored servo system. Positioning information is radially encoded in evenly-spaced servo bursts on each track. These servo burst wedges provide radial position information for each data head. Because the drive uses multiple zone recording, where each zone has a different bit density, split data fields are necessary to optimally utilize the non-servo area of the disk. The split data fields are achieved by special processing through the DCIIA, and their presence is transparent to the host system. The servo area remains phase coherent across the surface of the disk, even though the disk is divided into various data zones. The main advantage of the sectored servo systems is that the data heads are also servo heads, which means that sectored servo systems eliminate the problems of static and dynamic offsets between heads on different surfaces.

The Quantum Bigfoot CY 2.1/4.3/6.4AT hard disk drive servo system is also classified as a digital servo because track following compensation is done in firmware. The bump detect, on-track, velocity profiles, and other "housekeeping" tasks are also done in firmware.

Magneto-Resistive (MR) head technology is used to achieve high tracks per inch (TPI). MR technology employs separate read and write elements, which are not guaranteed to have colinear track centers with respect to a line drawn tangent to any of the concentric tracks. For this reason, the servo system is designed to track at different locations for read and write operations to maximize signal and minimize adjacent track interference, respectively.

The A, B, C, and D bursts are written in two passes per track, alternating between 2/3 and 1/3 track pitch steps. The bursts are not trimmed, which avoids the non-linearities associated with trimming due to gap length and writer throat height. Fine position is decoded by using the A-B phase around even track centers and C-D around odd track centers. With reference to even track center, B-C is used around the +50% offset and A-D is used around -50% offset. Two calibrated slopes are used, one for A-B and C-D and one for B-C and A-D.

5.3.2 Servo Burst and Track Information

Positional information is encoded on all tracks on all data surfaces. All data heads are also servo heads. The areas with servo/position information are called wedge areas. These wedge areas are evenly spaced radially around the disk, like spokes on a wheel.

Every wedge consists of five separate fields:

1. Automatic Gain Control (AGC)
2. Sync
3. Servo Address Mark (SAM)
4. Track Number
5. Burst Area

Because a phase-lock-loop is not used in the servo wedge area, time discrimination is used. Timing for all four fields is generated from the same crystal reference. The clock has a period of 25 ns, and will be referred to as T. Data is encoded in the following manner:

Servo Data Bit "0" = 10 000 010 0

Servo Data Bit "1" = 10 010 000 0

Because each digit has a width of T, the length of a servo bit is 9T.

5.4 READ AND WRITE OPERATIONS

The following paragraphs provide descriptions of the read channel, write channel, and IDE interface control operations.

5.4.1 The Read Channel

The drive has one read/write head for each data surface (two for the Quantum Bigfoot CY 2.1AT drive, four for the Quantum Bigfoot CY 4.3AT drive, and six for the Quantum Bigfoot CY 6.4AT drive). The signal path for the read channel begins at the read/write heads. As the magnetic flux transitions recorded on a disk pass under a head, they generate low-amplitude, differential output voltages. These read signals pass from the read/write head to the flex circuit's read preamplifier, which amplifies the signal. To ensure a high signal to noise ratio, preamplification occurs on the flex circuit because of its proximity to the heads.

The flex circuit transmits the preamplified signal from the HDA to the drive PCB. On the PCB, the Read/Write ASIC further processes the read signal to reduce ambiguities (for example, drop-ins, drop-outs, and peak shift). In addition, it converts the signal from the serial encoded head data to a synchronized data stream, with its accompanying clock. The Read/Write ASIC then sends the resynchronized and decoded data output to Quantum's proprietary Disk Controller and IDE Interface ASIC (DCIIA).

The DCIIA manages the flow of data between on the Read/Write ASIC and its IDE Interface Controller. It also controls data access for the external RAM buffer. The DCIIA format provides a serial bit stream. This NRZ (Non-Return to Zero) serial data is converted to an 8-bit byte. The Sequencer module identifies the data as belonging to the target sector. Data is presented to the host in a 16-bit word.

After a full sector is read, the DCIIA checks to see if the firmware needs to apply ECC on-the-fly or single-, double-, or triple-burst correction to the data. The buffer controller section of the DCIIA stores the data in the Cache and transmits the data to the IDE Interface Controller module, which transmits the data to the IDE bus.

5.4.2 The Write Channel

For the write channel, the signal path follows the reverse order of that for the read channel. The host presents a 16-bit word of data, by means of the IDE bus, to the DCIIA IDE Interface Controller. The Buffer Controller section of the DCIIA stores the data in the cache. Because data can be presented to the drive at a rate that exceeds the rate at which the drive can write data to a disk, data is stored temporarily in the cache. Thus, the host can present data to the drive at a rate that is independent of the rate at which the drive can write data to the disk.

Upon correct identification of the target address, the data is shifted to the Sequencer where an error correcting code is generated and appended. The Sequencer then converts

the bytes of data to a serial bit stream. The DCIIA transmits the data to the Read/Write ASIC where the data is encoded and precompensated to reduce intersymbol interference. The data is then transmitted to the Write Driver by means of the write data lines.

The drive's DCIIA switches the Preamplifier Write Driver IC to write mode and selects a head. Once the Write driver receives a write gate signal, it transmits current reversals to the head, which induces magnetic transitions on the disk.

5.4.3 Interface Control

The interface with the host system is through a 40-pin IDE interface connector. The DCIIA IDE Interface Controller module implements the IDE interface logic. Operating under the drive's µprocessor control, the DCIIA receives and transmits words of data over the IDE bus.

The DCIIA Buffer Controller writes data to or reads data from the Cache over 16 data lines. Under µController direction, the DCIIA controls the transfer of data and handles the addressing of the Cache. The internal data transfer rate to and from the Cache is 32 MB/s. This high transfer rate allows the DCIIA to communicate over the IDE interface at a PIO data transfer rate of 6.67 MB/s without using IORDY, up to 16.67 MB/s with PIO using IORDY, or a DMA transfer rate of up to 16.67 MB/s while it simultaneously controls disk-to-RAM transfers and microcontroller access to control code stored in the buffer RAM.

5.5 FIRMWARE FEATURES

This section describes the following firmware features:

- Disk caching
- Track and cylinder skewing
- Error detection and correction
- Defect management

5.5.1 Disk Caching

The Quantum Bigfoot CY 2.1/4.3/6.4AT hard disk drives incorporate DisCache, an 78 K disk cache, to enhance drive performance. This integrated feature is user-programmable and can significantly improve system throughput. Read and write caching can be enabled or disabled by using the Set Configuration or Set Features command.

5.5.1.1 Adaptive Caching

The cache buffer for the Quantum Bigfoot CY drives feature adaptive segmentation for more efficient use of the buffer's RAM. With this feature, the buffer space used for read and write operations is dynamically allocated. The cache can be flexibly divided into several segments, under program control. Each segment contains one cache entry.

A cache entry consists of the requested read data plus its corresponding prefetch data. Adaptive segmentation allows the drive to make optimum use of the buffer, the amount of stored data can be increased.

5.5.1.2 Read Cache

DisCache anticipates host-system requests for data and stores that data for faster access. When the host requests a particular segment of data, the caching feature uses a prefetch strategy to "look ahead" and automatically store the subsequent data from the disk into high-speed RAM. If the host requests this subsequent data, the RAM is accessed rather than the disk.

Since typically 50 percent or more of all disk requests are sequential, there is a high probability that subsequent data requested will be in the cache. This cached data can be retrieved in microseconds rather than milliseconds. As a result, DisCache can provide substantial time savings during at least half of all disk requests. In these instances, DisCache could save most of the disk transaction time by eliminating the seek and rotational latency delays that dominate the typical disk transaction. For example, in a 1K data transfer, these delays make up to 90 percent of the elapsed time.

DisCache works by continuing to fill its cache memory with adjacent data after transferring data requested by the host. Unlike a noncaching controller, Quantum's disk controller continues a read operation after the requested data has been transferred to the host system. This read operation terminates after a programmed amount of subsequent data has been read into the cache segment.

The cache memory consists of an 78 K DRAM buffer allocated to hold the data, which can be directly accessed by the host by means of the READ and WRITE commands. The memory functions as a group of segments (ring buffers) with rollover points at the end of each segment (buffer). The unit of data stored is the logical block (that is, a multiple of the 512 byte sector). Therefore, all accesses to the cache memory must be in multiples of the sector size.The following commands force emptying of the cache:

- RESET
- WRITE LONG
- EXECUTE DRIVE DIAGNOSTIC
- FORMAT TRACK
- READ CONFIGURATION
- READ DEFECT LIST
- READ LONG
- IDENTIFY DRIVE
- SET CONFIGURATION
- WRITE BUFFER
- SLEEP

5.5.1.3 Write Cache

When a write command is executed with write caching enabled, the drive stores the data to be written in a DRAM cache buffer and immediately sends a GOOD STATUS message to the host before the data is actually written to the disk. The host is then free to move on to other tasks, such as preparing data for the next data transfer, without having to wait for the drive to seek to the appropriate track or rotate to the specified sector.

While the host is preparing data for the next transfer, the drive immediately writes the cached data to the disk, usually completing the operation in less than 20 ms after issuing GOOD STATUS. With WriteCache, a single-block, random write, for example, requires only about 3 ms of host time. Without WriteCache, the same operation would occupy the host for about 20 ms.

WriteCache allows data to be transferred in a continuous flow to the drive rather than as individual blocks of data separated by disk access delays. This is achieved by taking advantage of the ability to write blocks of data sequentially on a disk that is formatted with a 1:1 interleave. This means that as the last byte of data is transferred out of the write cache and the head passes over the next sector of the disk, the first byte of the of the next block of data is ready to be transferred, thus there is no interruption or delay in the data transfer process.

The WriteCache algorithm fills the cache buffer with new data from the host while simultaneously transferring the data to the disk that the host previously stored in the cache.

5.5.1.4 Performance Benefits

In a drive without DisCache, there is a delay during sequential reads because of the rotational latency even if the disk actuator already is positioned at the desired cylinder. DisCache eliminates this rotational latency time (8.33 ms on average) when requested data resides in the cache.

Moreover, the disk must often service requests from multiple processes in a multitasking or multiuser environment. In these instances, while each process might request data sequentially, the disk drive must share time among all these processes. In most disk drives, the heads must move from one location to another. With DisCache, even if another process interrupts, the drive continues to access the data sequentially from its high-speed memory. In handling multiple processes, DisCache achieves its most impressive performance gains, saving both seek and latency time when desired data resides in the cache.

The cache can be flexibly divided into several segments, under program control. Each segment contains one cache entry. A cache entry consists of the requested read data plus its corresponding prefetch data.

The requested read data takes up a certain amount of space in the cache segment so the corresponding prefetch data can essentially occupy the rest of the space within the segment. The other factors determining prefetch size are the maximum and minimum prefetch. The drive's prefetch algorithm dynamically controls the actual prefetch value based on the current demand, with the consideration of overhead to subsequent commands.

5.5.2 Track and Cylinder Skewing

Track and cylinder skewing in the Quantum Bigfoot CY 2.1/4.3/6.4AT drives minimizes latency time and thus increases data throughput.

5.5.2.1 Track Skewing

Track skewing reduces the latency time that results when the drive must switch read/write heads to access sequential data. A track skew is employed such that the next logical sector of data to be accessed will be under the read/write head once the head switch is made and the data is ready to be accessed. Thus, when sequential data is on the same cylinder but on a different disk surface, a head switch is needed but not a seek. Since the sequential head-switch time is well defined on the Quantum Bigfoot CY 2.1/4.3/6.4AT drives, the sector addresses can be optimally positioned across track boundaries to minimize the latency time during a head switch. See Table 5-2.

5.5.2.2 Cylinder Skewing

Cylinder skewing also is used to minimize the latency time associated with a single-cylinder seek. The next logical sector of data that crosses a cylinder boundary is positioned on the drive such that after a single-cylinder seek is performed, and when the drive is ready to continue accessing data, the sector to be accessed is positioned directly under the read/write head. Therefore, the cylinder skew takes place between the last sector of data on the last head of a cylinder and the first sector of data on the first head of the next cylinder. Since the wedge-to-wedge time is constant over the entire disk, a single set of track and cylinder skew offsets will fulfill the requirement for all recording zones. See Table 5-2.

Table 5-2 *Track and Cylinder Skews*

Skew Type	Switch Time	Wedge Offset
Track Skew	3.57 ms	2.4 ms
Cylinder Skew	3.57 ms	2.4 ms

Note: 1. A wedge-to-wedge time of 148.44 ms is used, calculated using worst-case spindle variation (-0.25%) to provide a safety margin.

2. Wedge offsets are rounded up.

5.5.2.3 ID-less Format

The Quantum Bigfoot CY 2.1/4.3/6.4AT hard disk drives utilize an ID-less format. The ID-less format has several advantages over the traditional "ID After Wedge" or "ID Before Sector" methods of tracking the location of the actuator. For example, the lack of an ID field written on the disk gains approximately 4% of the overall track "real estate," which increases the total capacity. Secondly, since no IDs have to be read or corrected in case of an error, overall throughput is increased

In the ID-less environment, the track and cylinder skewing is based on units of wedges instead of the traditional sectors. The DCIIA ASIC contains a "Wedge Skew Register" to assist the task of skewing where the skew offset must now be calculated with every read/write operation. The firmware programs the skew offset into this register every time the drive goes to a new track. The DCIIA ASIC then adds this value to the wedge number in the sector descriptor, effectively relocating the "first sector" of the track away from the index. For example, if, without skew, sector 0 is to be found following wedge 0, then if the skew register is set to 10, sector will be found following wedge 10.

5.5.3 Error Detection and Correction

As disk drive areal densities increase, obtaining extremely low error rates requires a new generation of sophisticated error correction codes. Quantum Bigfoot CY 2.1/4.3/6.4AT hard disk drive series implement 224-bit triple-burst Reed-Solomon error correction techniques to reduce the uncorrectable read error rate to less than one bit in 1×10^{14} bits read.

When errors occur, an automatic retry, a double-burst, and a more rigorous triple-burst correction algorithm enable the correction of any sector with three bursts of four incorrect bytes each, or up to 12 multiple random one-byte burst errors. In addition to these advanced error correction capabilities, the drive uses an additional cross-checking code and algorithm, to double check the main ECC correction. This greatly reduces the probability of a miscorrection.

5.5.3.1 Background Information on Error Correction Code and ECC On-the-Fly

A sector on the Bigfoot CY 2.1/4.3/6.4AT drive is comprised of 512 bytes of user data, followed by four cross-checking (XC) bytes (32 bits), followed by 24 ECC check bytes (192 bits). The four cross-checking bytes are used to double check the main ECC correction and reduce the probability of miscorrection. Errors of up to 64 bits within one sector can be corrected "on-the-fly," in real time as they occur, allowing a high degree of data integrity with no impact on the drive's performance.

The drive does not need to re-read a sector on the next disk revolution or apply ECC for those errors that are corrected on-the-fly. Errors corrected in this manner are invisible to the host system.

When errors cannot be corrected on-the-fly, an automatic retry, and a more rigorous triple-burst error correction algorithm enables the correction of any sector with three bursts of four incorrect bytes each (up to 12 contiguous bytes), or up to 12 multiple random one-byte burst errors. In addition to this error correction capability, the drive's implementation of an additional cross-checking code and algorithm double checks the main ECC correction, and greatly decreases the likelihood of miscorrection.

The 24 ECC check bytes shown in Figure 5-6 are used to detect and correct errors. The cross-checking and ECC data is computed and appended to the user data when the sector is first written.

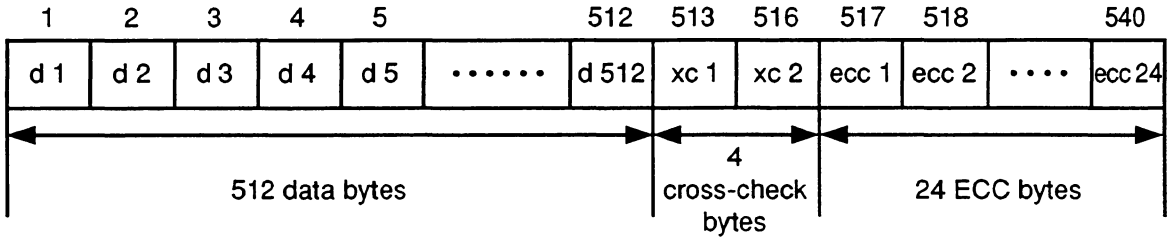

Figure 5-6 *Sector Data Field with ECC Check Bytes*

To obtain the ECC check byte values, each byte (including cross-checking and ECC bytes) within the sector is interleaved into one of three groups, where the first byte is in interleave 1, the second byte is in interleave 2, the third byte is in interleave 3, the fourth byte is in interleave 1, the fifth byte is in interleave 2, and so on, as shown in Figure 5-7.

Interleave 1 ⇨	d1	d5	• • • •	d508	d512	xc4	ecc4	ecc8	ecc12	ecc16	ecc20	ecc24
Interleave 2 ⇨	d2	d6	• • • •	d509	xc1	ecc1	ecc5	ecc9	ecc13	ecc17	ecc21	
Interleave 3 ⇨	d3	d7	• • • •	d510	xc2	ecc2	ecc6	ecc10	ecc14	ecc18	ecc22	
Interleave 4 ⇨	d4	d8	• • • •	d511	xc3	ecc3	ecc7	ecc11	ecc15	ecc19	ecc23	

Figure 5-7 *Byte Interleaving*

Note: ECC interleaving is not the same as the sector interleaving that is done on the disk.

Each of the four interleaves is encoded with six ECC bytes, resulting in the 24 ECC bytes at the end of the sector. The four cross checking bytes are derived from all 512 data bytes.

The combination of the interleaving, and the nature of the ECC formulas enable the drive to know where the error occurs.

Because the ECC check bytes follow the cross checking bytes, errors found within the cross-checking bytes can be corrected. Due to the power and sophistication of the code, errors found within the ECC check bytes can also be corrected.

Each time a sector of data is read, the Bigfoot CY 2.1/4.3/6.4AT drives will generate a new set of ECC check bytes and cross-checking bytes from the user data. These new check bytes are compared to the ones originally written to the disk. The difference between the newly computed and original check bytes is reflected in a set of 24 *syndromes* and four cross checking syndromes, which correspond to the number of check bytes. If all the syndrome values equal zero, the data was read with no errors, and the sector is transferred to the host system. If any of the syndromes do not equal zero, an error has occurred. The type of correction the drive applies depends on the nature and the extent of the error.

High speed on-the-fly error correction saves several milliseconds on each single-, or double- burst error, because there is no need to wait for a disk revolution to bring the sector under the head for re-reading.

Correction of Single-, or Double-Burst Errors On-the-Fly

Single-burst errors may have up to four erroneous bytes (32 bits) within a sector, provided that each of the four bytes occur in a different interleave.

The Bigfoot CY 2.1/4.3/6.4AT drives have the capability to correct double-burst errors on-the-fly as well. Double-burst errors can be simply viewed as two spans of errors within one sector. More specifically, correctable double-burst errors must have two or fewer erroneous bytes per interleave.

The drive's Reed-Solomon ECC corrects double-burst errors up to 64 bits long, (provided that the error consists of two or fewer bytes residing in each of the interleaves).

Double-Burst Error Examples

In the example shown in Figure 5-8 C, the 58-bit error is uncorrectable since it occupies more than two erroneous bytes per interleave.

The other two 64-bit errors, shown in Figure 5-8 A and B, are correctable because no more than two error bytes of the entire error reside in any one of the interleaves.

Note: Any 57-bit error burst can be corrected on-the-fly using double-burst error correction because no more than two bytes can occupy each interleave.

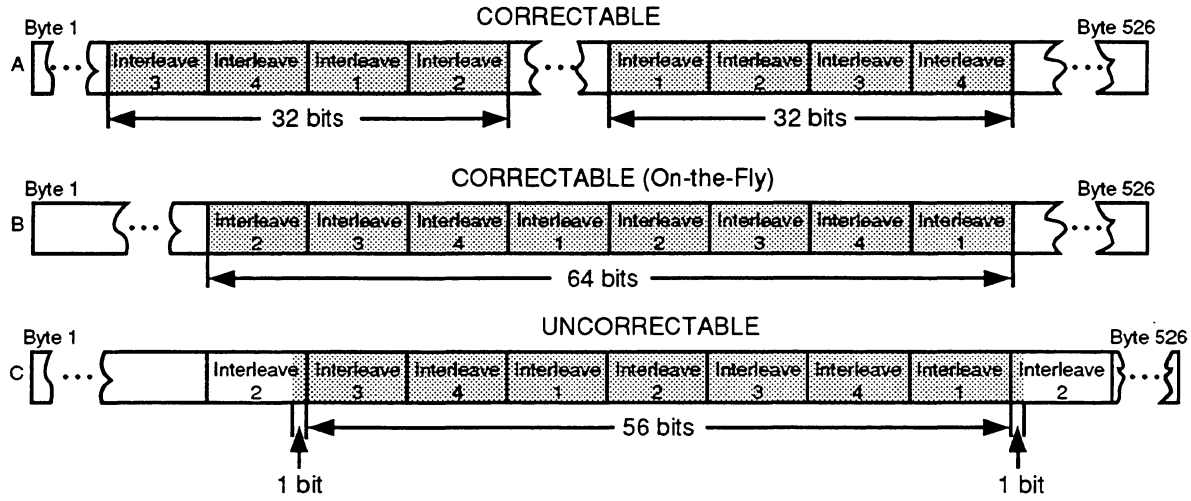

Figure 5-8 *Correctable and Uncorrectable Double-Burst Errors*

Correction of Triple-Burst Errors

Through sophisticated algorithms, Bigfoot CY 2.1/4.3/6.4AT drives have the capability to correct triple-burst errors, even though the probability of their occurrence is low. Triple-burst errors can be simply viewed as three spans of errors within one sector. More specifically, correctable triple-burst errors must have three or fewer erroneous bytes per interleave, and will not be corrected on-the-fly.

The drive's Reed-Solomon ECC corrects triple-burst errors up to 96 bits long, (provided that the error consists of three or fewer bytes residing in each of the interleaves).

If the triple-burst correction is successful, the data from the sector can be written to a spare sector, and the logical address will be mapped to the new physical location.

Triple-Burst Error Examples

In the example shown in Figure 5-9 C, the 90-bit error is uncorrectable since it occupies more than three erroneous bytes per interleave.

The other two 96-bit errors, shown in Figure 5-9 A and B, are correctable because no more than three error bytes of the entire error reside in any one of the interleaves.

Note: Any 89-bit error burst can be corrected using triple-burst error correction because no more than three bytes can occupy each interleave.

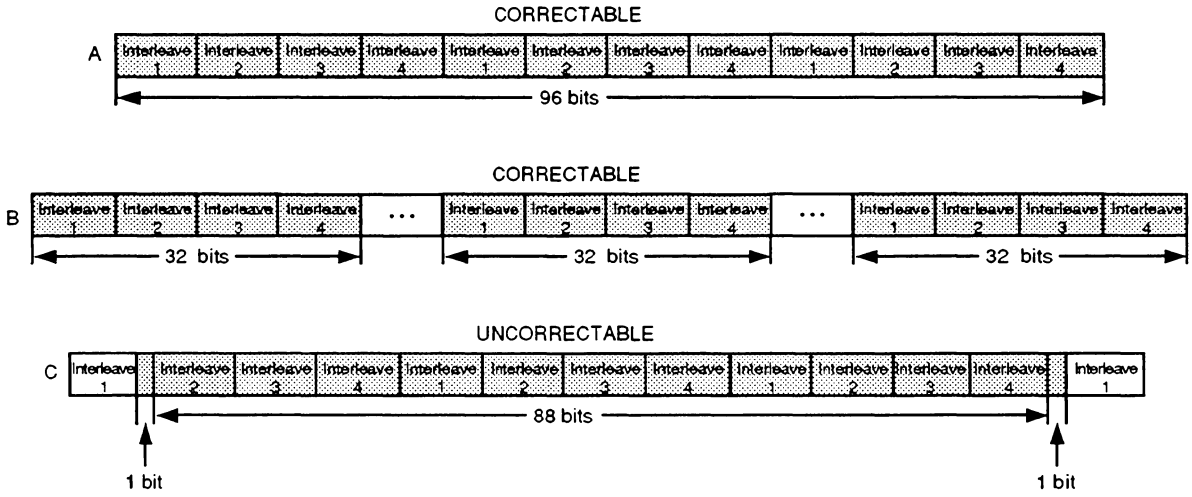

Figure 5-9 *Correctable and Uncorrectable Triple-Burst Errors*

Multiple Random Burst Errors

The drive's ECC can correct up to 96 bits of multiple random errors, provided that the incorrect bytes follow the guidelines for correctable triple-burst errors. Up to 64 bits of multiple random errors can be corrected on-the-fly, provided that the incorrect bytes follow the guidelines for correctable double-burst errors. Up to 24 bits of multiple random errors can be corrected on-the-fly if two bytes per interleave contains an error. If more than three bytes in any one interleave are in error, the sector cannot be corrected. Figure 5-10 shows an example of a correctable random burst error consisting of 12 bytes (96 bits). This random burst error is correctable because no more than three bytes within each interleave are in error.

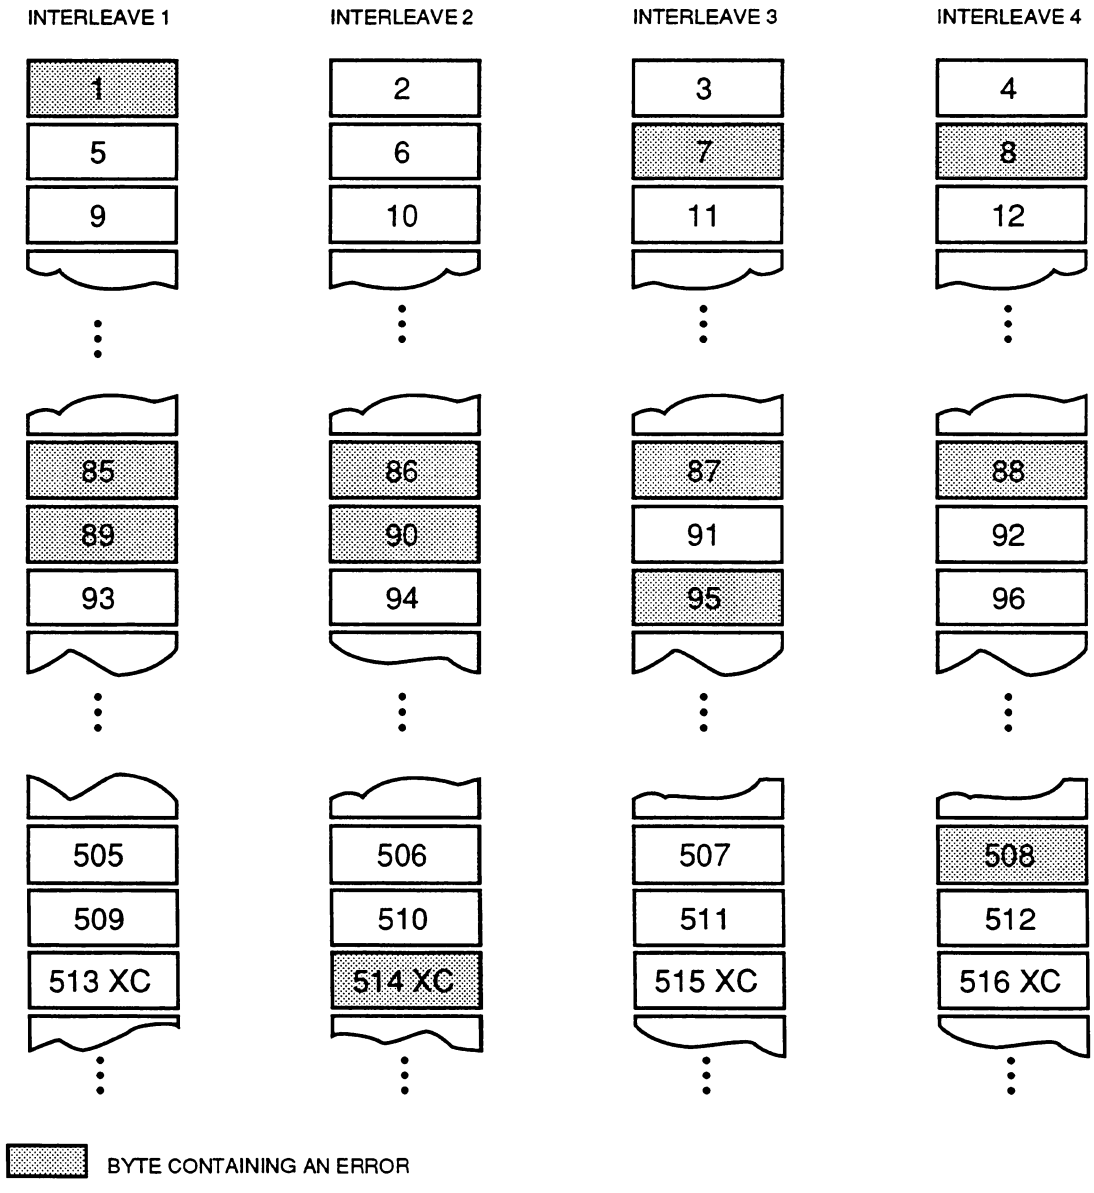

Figure 5-10 *Nine Correctable Random Burst Errors*

5.5.3.2 ECC Error Handling

When a data error occurs, the Quantum Bigfoot CY 2.1/4.3/6.4AT hard disk drives check to see if the error is correctable on-the-fly. This process takes about 200 µs. If the error is correctable on-the-fly, the error is corrected and the data is transferred to the host system.

If the data is not correctable on-the-fly, the sector is re-read in an attempt to read the data correctly without applying the triple-burst ECC correction. Before invoking the complex triple-burst ECC algorithm, the drive will always try to recover from an error by attempting to re-read the data correctly. This strategy prevents invoking correction on non-repeatable errors. Each time a sector in error is re-read a set of ECC syndromes is computed. If all of the syndrome values equal zero, the data was read with no errors, and the sector is transferred to the host system. If any of the syndrome values do not equal zero, an error has occurred, the syndrome values are retained, and another re-read is invoked.

Note: Non-repeatable errors are usually related to the signal to noise ratio of the system. They are not due to media defects.

When the sets of syndromes from two consecutive re-reads are the same, a stable syndrome has been achieved. This event may be significant depending on whether the automatic read reallocation or early correction features have been enabled. If the early correction feature has been enabled and a stable syndrome has been achieved, triple-burst ECC correction is applied, and the appropriate message is transferred to the host system (e.g., corrected data, etc.).

Note: These features can be enabled or disabled through the ATA Set Configuration command. The EEC bit enables early ECC triple-burst correction if a stable syndrome has been achieved before all of the re-reads have been exhausted. The ARR bit enables the automatic reallocation of defective sectors.

If the automatic read reallocation feature is enabled, the drive, when encountering triple-burst errors, will attempt to re-read up to 8 times the retry count set in the AT Configuration bytes.

Note: The Quantum Bigfoot CY 2.1/4.3/6.4AT drives are shipped from the factory with the automatic read reallocation feature enabled so that any new defective sectors can be easily and automatically reallocated for the average AT end user.

5.5.4 Defect Management

The Quantum Bigfoot CY 2.1/4.3/6.4AT drives allocate thousands of sectors at the end of the physical drive to accommodate inline sparing of defective sectors. In the factory, the media is scanned for defects. If a sector is found to be defective, the address of the sector is added to the drive's defect list. Sectors located physically subsequent to the defective sector are assigned logical block addresses such that a sequential ordering of logical blocks is maintained. The inline sparing technique is maintained in an attempt to eliminate slow data transfer that would result from a single defective sector on a cylinder. All factory defective sectors are inline spared.

Defects that occur in the field are known as *grown* defects. If such a defective sector is found in the field, inline sparing is not performed on these grown defects. Instead, the sector is reallocated to an available spare sector on a nearby cylinder.

Sectors are considered to contain grown defects if the triple-burst ECC algorithm must be applied to recover the data. If this algorithm is successful, the corrected data is stored in the newly allocated sector. If the algorithm is not successful, the erroneous data is stored in the newly allocated sector, and a flag is set in the data ID field that causes the drive to report an ECC error each time the sector is read. This condition remains until the sector is rewritten.

Chapter 6
IDE BUS INTERFACE AND ATA COMMANDS

This chapter describes the interface between Quantum Bigfoot CY 2.1/4.3/6.4AT hard disk drives and the IDE bus. The commands that are issued from the host to control the drive are listed, as well as the electrical and mechanical characteristics of the interface.

6.1 INTRODUCTION

Quantum Bigfoot CY 2.1/4.3/6.4 AT hard disk drives use the standard IBM PC IDE bus interface, and are compatible with systems that provide an IDE interface connector on the motherboard. It may also be used with a third-party adapter board in systems that do not have a built-in IDE adapter. The adapter board plugs into a standard 16-bit expansion slot in an AT-compatible computer. A cable connects the drive to the adapter board.

6.2 SOFTWARE INTERFACE

The Quantum Bigfoot CY 2.1/4.3/6.4 AT drives are controlled by the Basic Input/Output System (BIOS) program residing in an IBM PC AT, or compatible PC. The BIOS communicates directly with the drive's built-in controller. It issues commands to the drive and receives status information from the drive.

6.3 MECHANICAL DESCRIPTION

6.3.1 Drive Cable and Connector

The hard disk drive connects to the host computer by means of a cable. This cable has a 40-pin connector that plugs into the drive, and a 40-pin connector that plugs into the host computer. At the host end, the cable plugs into either an adapter board residing in a host expansion slot or an on-board IDE adapter.

If two drives are connected by a cable with two 40-pin drive connectors, the cable-select feature of the Quantum Bigfoot CY 2.1/4.3/6.4 AT drives automatically configure each as either drive 0 or drive 1 depending on the configuration of pin 28 on the connector. See Section 3.3.1, "Cable Select/Slave Present (CS/SP) Jumper," for more information about the cable select jumper.

6.4 ELECTRICAL INTERFACE

6.4.1 IDE Bus Interface

A 40-pin IDE interface connector on the motherboard or an adapter board provides an interface between the drive and a host that uses an IBM PC AT bus. The IDE interface contains bus drivers and receivers compatible with the standard AT bus. The AT-bus interface signals D8–D15, INTRQ, and IOCS16– require the IDE adapter board to have an extended I/O-bus connector.

The IDE interface buffers data and control signals between the drive and the AT bus of the host system, and decodes addresses on the host address bus. The Command Block Registers on the drive accept commands from the host system BIOS.

Note: Some host systems do not read the Status Register after the drive issues an interrupt. In such cases, the interrupt may not be acknowledged. To overcome this problem, you may have to configure a jumper on the motherboard or adapter board to allow interrupts to be controlled by the drive's interrupt logic. Read your motherboard or adapter board manual carefully to find out how to do this.

6.4.1.1 Electrical Characteristics

All signals are transistor-transistor logic (TTL) compatible–with logic 1 greater than 2.0 volts and less than 5.25 volts; and logic 0 greater than 0.0 volts and less than 0.8 volts. Neither the adapter board, motherboard interface, or drives require terminating resistors.

6.4.1.2 Drive Signals

The drive connector (J1, section C) connects the drive to an adapter board or onboard IDE adapter in the host computer. J1, section C is a 40-pin shrouded connector with two rows of 20 pins on 100-mil centers. J1 has been keyed by removing pin 20. The connecting cable is a 40-conductor flat ribbon cable, with a maximum length of 18 inches.

Table 6-1 describes the signals on the drive connector (J1, section C). The drive does not use all of the signals provided by the IDE bus. Table 6-2 shows the relationship between the drive connector (J1, section C) and the IDE bus.

Note: In Table 6-1 , the following conventions apply:

A minus sign follows the name of any signal that is asserted as active low.
Direction (DIR) is in reference to the drive.
IN indicates input to the drive.
OUT indicates output from the drive.
I/O indicates that the signal is bidirectional.

Table 6-1 *Drive Connector Pin Assignments (J1, Section C)*

SIGNAL	NAME	DIR	PIN	DESCRIPTION
Reset	RESET-	IN	1	Drive reset signal from the host system, inverted on the adapter board or motherboard. Asserted for at least 300 ns during start up and deasserted thereafter, unless some event subsequently requires that the drive be reset.
Ground	Ground	–	2	Ground between the host system and the drive.
Data Bus		I/O	3-18	An 8/16-bit, bidirectional data bus between the host and the drive. D0-D7 are used for 8-bit transfers, such as registers and ECC bytes.
	DD0		17	Bit 0
	DD1		15	Bit 1
	DD2		13	Bit 2
	DD3		11	Bit 3
	DD4		9	Bit 4
	DD5		7	Bit 5
	DD6		5	Bit 6
	DD7		3	Bit 7
	DD8		4	Bit 8
	DD9		6	Bit 9
	DD10		8	Bit 10
	DD11		10	Bit 11
	DD12		12	Bit 12
	DD13		14	Bit 13
	DD14		16	Bit 14
	DD15		18	Bit 15
Ground	Ground	–	19	Ground between the host system and the drive.
Keypin	KEYPIN	–	20	Pin removed to key the interface connector.
DMA Request	DMARQ	OUT	21	Asserted by the drive when it is ready to exchange data with the host. The direction of the data transfer is determined by DIOW- and DIOR-. DMARQ is used in conjunction with DMACK-
Ground	Ground	–	22	Ground between the host system and the drive.
I/O Write	DIOW-	IN	23	The rising edge of this write strobe provides a clock for data transfers from the host data bus (DD0-DD7 or DD0-DD15) to a register or to the drive's data port.
Ground	Ground	–	24	Ground between the host system and the drive.
I/O Read	DIOR-	IN	25	The rising edge of this read strobe provides a clock for data transfers from a register or the drive's data port to the host data bus (DD0-DD7 or DD0-DD15). The rising edge of DIOR- latches data at the host.

SIGNAL	NAME	DIR	PIN	DESCRIPTION
Ground	Ground	–	26	Ground between the host system and the drive.
I/O Channel Ready	IORDY	OUT	27	When the drive is not ready to respond to a data transfer request, the IORDY signal is asserted active low to extend the host transfer cycle of any host register read or write access. When IORDY is deasserted, it is in a high-impedance state and it is the host's responsibility to pull this signal up to a high level (if necessary).
Cable Select		–	28	A signal from the host that allows the drive to be configured as drive 0 when the signal is 0 (grounded), and as drive 1 when the signal is 1 (high).
DMA Acknowledge	DACK1–	IN	29	Used by the host to respond to the drive's DMARQ signal. DMARQ signals that there is more data available for the host.
Ground	Ground	–	30	Ground between the host system and the drive.
Interrupt Request	INTRQ	OUT	31	An interrupt to the host system. Asserted only when the drive microprocessor has a pending interrupt, the drive is selected, and the host clears nIEN in the Device Control Register. When nIEN is a 1 or the drive is not selected, this output signal is in a high-impedance state, regardless of the presence or absence of a pending interrupt. INTRQ is deasserted by an assertion of RESET–, the setting of SRST in the Device Control Register, or when the host writes to the Command Register or reads the Status Register. When data is being transferred in programmed I/O (PIO) mode, INTRQ is asserted at the beginning of each data block transfer. Exception: INTRQ is not asserted at the beginning of the first data block transfer that occurs when any of the following commands executes: FORMAT TRACK, Write Sector, WRITE BUFFER, or WRITE LONG.
16-Bit I/O	IOCS16–	OUT	32	An open-collector output signal. Indicates to the host system that the 16-bit data port has been addressed and that the drive is ready to send or receive a 16-bit word. When transferring data in PIO mode, if IOCS16– is not asserted, D0–D7 are used for 8-bit transfers; if IOCS16– is asserted, D0–D15 are used for 16-bit data transfers.
Drive Address Bus				A 3-bit, binary-coded address supplied by the host when accessing a register or the drive's data port.
Bit 1	DA1	IN	33	
Bit 0	DA0	IN	35	
Bit 2	DA2	IN	36	

SIGNAL	NAME	DIR	PIN	DESCRIPTION
Passed Diagnostics	PDIAG–	I/O	34	Drive 0 (Master) monitors this Drive 1 (Slave) open-collector output signal, which indicates the result of a diagnostics command or reset. Each drive has a 10K pull-up resistor on this signal. Following the receipt of a power-on reset, software reset, or RESET- drive 1 negates PDIAG- within 1 ms. PDIAG- indicates to drive 0 that drive 1 is busy (BSY=1). Then, drive 1 asserts PDIAG- within 30 seconds, indicating that drive 1 is no longer busy (BSY=0) and can provide status information. Following the assertion of PDIAG–, drive 1 is unable to accept commands until drive 1 is ready (DRDY=1)–that is, until the reset procedure for drive 1 is complete. Following the receipt of a valid EXECUTE DRIVE DIAGNOSTIC command, drive 1 negates PDIAG- within 1 ms, indicating to drive 0 that it is busy and has not yet passed its internal diagnostics. If drive 1 is present, drive 0 waits for drive 1 to assert PDIAG- for up to 5 seconds after the receipt of a valid EXECUTE DRIVE DIAGNOSTIC command. Since PDIAG- indicates that drive 1 has passed its internal diagnostics and is ready to provide status, drive 1 clears BSY prior to asserting PDIAG–. If drive 1 fails to respond during reset initialization, drive 0 reports its own status after completing its internal diagnostics. Drive 0 is unable to accept commands until drive 0 is ready (DRDY=1)–that is, until the reset procedure for drive 0 is complete.
Chip Select 0	CS1FX–	IN	37	Chip-select signal decoded from the host address bus. Used to select the host-accessible Command Block Registers.
Chip Select 1	CS3FX–	IN	38	Chip select signal decoded from the host address bus. Used to select the host-accessible Control Block Registers.

SIGNAL	NAME	DIR	PIN	DESCRIPTION
Drive Active/Slave Present	DASP-	I/O	39	A time-multiplexed signal that indicates either drive activity or that drive 1 is present. During power-on initialization, DASP- is asserted by drive 1 within 400 ms to indicate that drive 1 is present. If drive 1 is not present, drive 0 asserts DASP- after 450 ms to light the drive-activity LED. An open-collector output signal, DASP- is deasserted following the receipt of a valid command by drive 1 or after the drive is ready, whichever occurs first. Once DASP- is deasserted, either hard drive can assert DASP- to light the drive-activity LED. Each drive has a 10K pull-up resistor on this signal. If an external drive-activity LED is used to monitor this signal, an external resistor must be connected in series between the signal and a +5 volt supply in order to limit the current to 24 mA maximum.
Ground	Ground	–	40	Ground between the host system and the drive.

6.4.1.3 IDE Bus Signals

See Table 6-2 for the relationship between the drive signals and the IDE bus.

Table 6-2 *Relationship of Drive Signals to the IDE Bus*

DRIVE CONNECTOR (J1)		DIRECTION	IDE BUS	
1	RESET-	<-(INV)	B2	RESET DRV
2	GROUND	–	–	GROUND
3	DD7	<->	A2	SD7
4	DD8	<->	C11	SD8
5	DD6	<->	A3	SD6
6	DD9	<->	C12	D9
7	DD5	<->	A4	SD5
8	DD10	<->	C13	SD10
9	DD4	<->	A5	SD4
10	DD11	<->	C14	SD11
11	DD3	<->	A6	SD3
12	DD12	<->	C15	SD12
13	DD2	<->	A7	SD2
14	DD13	<->	C16	SD13
15	DD1	<->	A8	SD1
16	DD14	<->	C17	SD14
17	DD0	<->	A9	SD0
18	DD15	<->	C18	SD15
19	GROUND	–	–	GROUND
20	KEYPIN	–	–	NO CONNECTION
21	DMARQ	->	–[1]	DRQ
22	GROUND	–	–	GROUND
23	DIOW-	<-	B13	DIOW-
24	GROUND	–	–	GROUND
25	DIOR-	<-	B14	IOR-
26	GROUND	–	–	GROUND
27	IORDY	->	A10	I/O CH RDY
28	CABLE SELECT	<-	–	NO CONNECTION
29	DMACK-	<-	–[2]	DACK-
30	GROUND	–	–	GROUND
31	INTRQ	->	D7	INTRQ
32	IOCS16-	->	D2	I/O CS16-
33	ADDR1	<-	A30	SA1
34	PDIAG-	–	–	NO CONNECTION
35	DA0	<-	A31	SA0
36	ADDR2	<-	A29	SA2
37	CS1FX-	<-	–[3]	CS0-
38	CS3FX-	<-	–[3]	CS1-
39	DASP-	–	–[3]	NO CONNECTION
40	GROUND	–	–	GROUND

1. DMARQ from the drive must be jumpered to one of the DRQ lines on the motherboard or host adapter (normally connected to DRQ6).
2. DACK- from the drive must be jumpered to one of the DACK- lines on the motherboard or host adapter (normally connected to DACK6-).
3. CS1FX-, CS3FX-, and DASP- originate on the adapter board.

6.4.2 Host Interface Timing

6.4.2.1 Programmed I/O (PIO) Transfer Mode

The PIO host interface timing shown in Table 6-3 is in reference to signals at 0.8 volts and 2.0 volts. All times are in nanoseconds, unless otherwise noted. Figure 6-1 provides a timing diagram.

Table 6-3 *PIO Host Interface Timing*

SYMBOL	DESCRIPTION	MIN/MAX	MODE 4[1] (local bus)	QUANTUM Bigfoot CY 2.1/4.3/6.4 AT
t0	Cycle Time	min	120	120
t1	Address Valid to DIOW-/DIOR-Setup	min	25	25
t2	DIOW-/DIOR- Pulsewidth (8- or 16-bit)	min	70	70
t2i	DIOW-/DIOR- Negated Pulsewidth	min	25	25
t3	DIOW-Data Setup	min	20	20
t4	DIOW- Data Hold	min	10	10
t5	DIOR- Data Setup	min	20	20
t5a	DIOR- to Data Valid	max	–	–
t6	DIOR- Data Hold	min	5	5
t6z	DIOR- Data Tristate	max	30	30
t7	Address Valid to IOCS16- Assertion	max	N/A	N/A
t8	Address Valid to IOSC16- Deassertion	max	N/A	N/A
t9	DIOW-/DIOR- to Address Valid Hold	min	10	10
tA	IORDY[2] Setup Time	min	35	35
tB	IORDY Pulse Width	max	1250	1250
tR	Read Data Valid to IORDY active (if IORDY is initially low after tA)	min	0	0

1. ATA Mode 4 timing is listed for reference only.
2. Transfer rates above 6 MB/s require the use of IORDY.

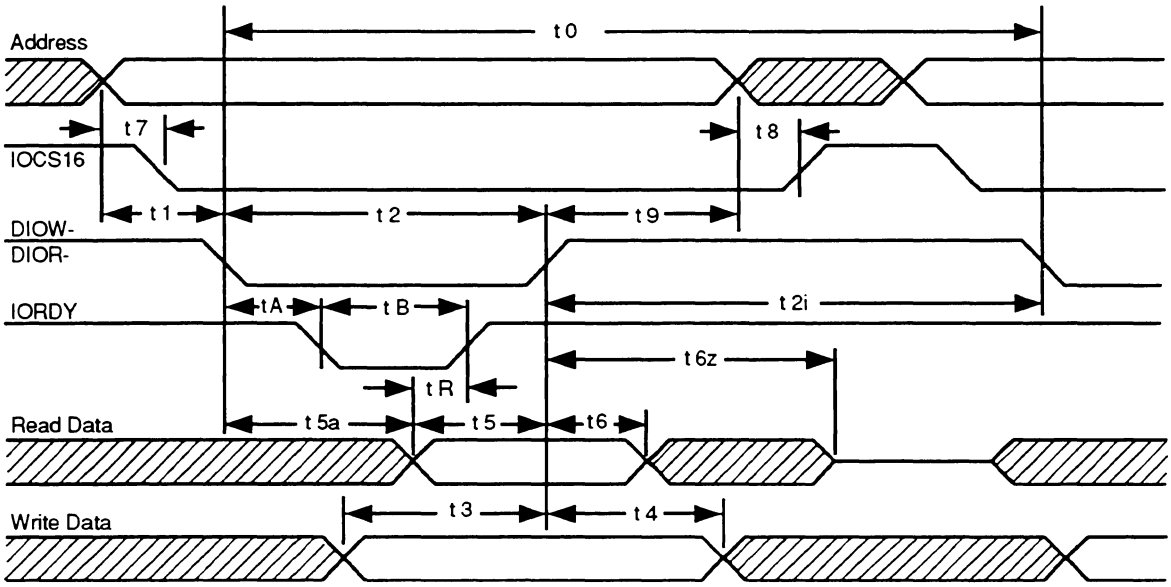

Figure 6-1 *PIO Interface Timing*

6.4.2.2 Multiword DMA Transfer Mode

The multiword DMA host interface timing shown in Table 6-4 is in reference to signals at 0.8 volts and 2.0 volts. All times are in nanoseconds, unless otherwise noted. Figure 6-2 provides a timing diagram.

Table 6-4 *Multiword DMA Host Interface Timing*

SYMBOL	DESCRIPTION	MIN/MAX	MODE 2[1] (local bus)	QUANTUM Bigfoot AT
t0	Cycle Time	min	120	120
tD	DIOR-/DIOW- Pulsewidth	min	70	70
tE	DIOR- to Data Valid	max	-	-
tF	DIOR- Data Hold	min	5	5
tFz	DIOR- Data Tristate[2]	max	20	20
tG	DIOW- Data Setup	min	20	20
tH	DIOW- Data Hold	min	10	10
tI	DMACK to DIOR-/DIOW- Setup	min	0	0
tJ	DIOR-/DIOW- to DMACK- Hold	min	5	5
tK	DIOR-/DIOW- Negated Pulsewidth	min	25	25
tL	DIOR-/DIOW- to DMARQ Delay	max	35	35
tz	DMACK- Data Tristate[3]	max	25	25

1. ATA Mode 2 timing is listed for reference only.
2. The Quantum Bigfoot CY 2.1/4.3/6.4 AT drive tristates after each word transferred.
3. Symbol tz only applies on the last tristate at the end of a multiword DMA transfer cycle.

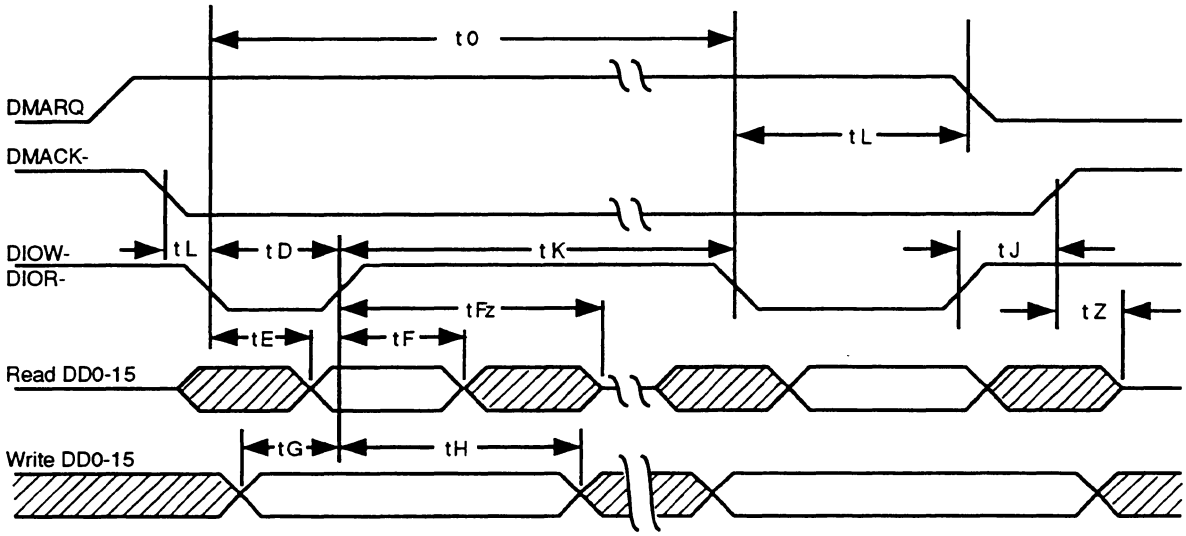

Figure 6-2 *Multiword DMA Bus Interface Timing*

6.4.2.3 Host Interface RESET Timing

The host interface RESET timing shown in Table 6-5 is in reference to signals at 0.8 volts and 2.0 volts. All times are in nanoseconds, unless otherwise noted. Figure 6-3 provides a timing diagram.

Table 6-5 *Host Interface RESET Timing*

SYMBOL	DESCRIPTION	MINIMUM	MAXIMUM
tM	RESET- Pulse width	300	–

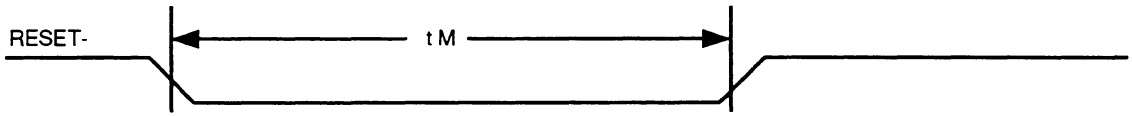

Figure 6-3 *Host Interface RESET Timing*

6.5 REGISTER ADDRESS DECODING

The host addresses the drive by using programmed I/O. Host address lines A0–A2, chip-select CS1FX– and CS3FX–, and IOR– and IOW– address the disk registers. Host address lines A3–A9 generate the two chip-select signals, CS1FX– and CS3FX–.

- Chip Select CS1FX– accesses the eight Command Block Registers.
- Chip Select CS3FX– is valid during 8-bit transfers to or from the Alternate Status Register.

The drive selects the primary or secondary command block addresses by setting Address bit A7.

Data bus lines 8–15 are valid only when IOCS16– is active and the drive is transferring data. The drive transfers ECC information only on data bus lines 0–7. Data bus lines 8–15 are invalid during transfers of ECC information.

I/O to or from the drive occurs over an I/O port that routes input or output data to or from selected registers, by using the following encoded signals from the host:

CS1FX–, CS3FX–, DA2, DA1, DA0, DIOR–, and DIOW–.

The host writes to the Command Block Registers when transmitting commands to the drive, and to the Control Block Registers when transmitting control, like a software reset. Table 6-6 lists the selection addresses for these registers.

Table 6-6 *I/O Port Functions and Selection Addresses*

FUNCTION		HOST SIGNALS				
CONTROL BLOCK REGISTERS		**CS1FX–**	**CS3FX–**	**DA2**	**DA1**	**DA0**
READ (DIOR–)	**WRITE (DIOW–)**					
Data Bus High Impedance	Not Used	N[1]	N	X[2]	X	X
Data Bus High Impedance	Not Used	N	A[3]	0	X	X
Data Bus High Impedance	Not Used	N	A	1	0	X
Alternate Status	Device Control	N	A	1	1	0
Drive Address	Not Used	N	A	1	1	1
COMMAND BLOCK REGISTERS						
READ (DIOR–)	**WRITE (DIOW–)**					
Data Port	Data Port	A	N	0	0	0
Error Register	Features	A	N	0	0	1
Sector Count	Sector Count	A	N	0	1	0
Sector Number[4]	Sector Number	A	N	0	1	1
LBA Bits 0–7[5]	LBA Bits 0–7	A	N	0	1	1
Cylinder Low[4]	Cylinder Low	A	N	1	0	0
LBA Bits 8–15[5]	LBA Bits 8–15	A	N	1	0	0
Cylinder High[4]	Cylinder High	A	N	1	0	1
LBA Bits 16–23[5]	LBA Bits 16–23	A	N	1	0	1
Drive/Head[4]	Drive/Head	A	N	1	1	0
LBA Bits 24–27[5]	LBA Bits 24–27	A	N	1	1	0
Status	Command	A	N	1	1	1
Invalid Address	Invalid Address	A	A	X	X	X

1. N = signal deasserted
2. X = signal either asserted or deasserted
3. A = signal asserted
4. Mapping of registers in CHS mode
5. Mapping of registers in LBA mode

After power on or following a reset, the drive initializes the Command Block Registers to the values shown in Table 6-7.

Table 6-7 *Command Block Register Initial Values*

REGISTER	VALUE
Error Register	01
Sector Count Register	01
Sector Number Register	01
Cylinder Low Register	00
Cylinder High Register	00
Drive/Head Register	00

6.6 REGISTER DESCRIPTIONS

The Quantum Bigfoot CY 2.1/4.3/6.4 AT hard disk drives emulate the ATA Command and Control Block Registers. Functional descriptions of these registers are given in the next two sections.

6.6.1 Control Block Registers

6.6.1.1 Alternate Status Register

The Alternate Status Register contains the same information as the Status Register in the command block. Reading the Alternate Status Register does not imply the acknowledgment of an interrupt by the host or clear a pending interrupt. See the description of the Status Register in section 6.6.2.8 for definitions of bits in this register.

6.6.1.2 Device Control Register

This write-only register contains two control bits, as shown in Table 6-8.

Table 6-8 *Device Control Register Bits*

BIT	MNEMONIC	DESCRIPTION
7	Reserved	-
6	Reserved	-
5	Reserved	-
4	Reserved	-
3	1	Always 1
2	SRST[1]	Host software reset bit
1	nIEN[2]	Drive interrupt enable bit
0	0	Always 0

1. SRST = Host Software Reset bit. When the host sets this bit, the drive is reset. When two drives are daisy-chained on the interface, this bit resets both drives simultaneously.
2. nIEN = Drive Interrupt Enable bit. When nIEN equals 0 or the host has selected the drive, the drive enables the host interrupt signal INTRQ through a tristate buffer to the host. When nIEN equals 1 or the drive is not selected, the host interrupt signal INTRQ is in a high-impedance state, regardless of the presence or absence of a pending interrupt.

6.6.1.3 Drive Address Register

The Drive Address Register returns the head-select addresses for the drive currently selected. Table 6-9 shows the Drive Address bits.

Table 6-9 *Drive Address Register Bits*

BIT	MNEMONIC	DESCRIPTION
7	**HiZ[1]**	**High Impedance bit**
6	nWTG[2]	Write Gate bit
5	nHS3[3]	Head Address msb
4	nHS2	-
3	nHS1	-
2	nHS0	Head Address lsb
1	nDS1[4]	Drive 1 Select bit
0	nDS0	Drive 0 Select bit

1. HiZ = High Impedance bit. When the host reads the register, this bit will be in a high impedance state.
2. nWTG = Write Gate bit. When a write operation to the drive is in progress, nWTG equals 0.
3. nHS0-nHS3 = Head Address bits. These bits are equivalent to the one's complement of the binary-coded address of the head currently selected.
4. nDS0-nDS1 = Drive Select bits. When drive 1 is selected, nDS1 equals 0. When drive 0 is selected, nDS0 equals 0.

6.6.2 Command Block Registers

6.6.2.1 Data Port Register

All data transferred between the device data buffer and the host passes through the Data Port Register. The host transfers the sector table to this register during execution of the FORMAT TRACK command. Transfers of ECC bytes during the execution of READ LONG or WRITE LONG commands are 8-bit transfers.

6.6.2.2 Error Register

The Error Register contains status information about the last command executed by the drive. The contents of this register are valid only when the Error bit (ERR) in the Status Register is set to 1. The contents of the Error Register are also valid at power on, and at the completion of the drive's internal diagnostics, when the register contains a status code. When the error bit in the Status Register is set to 1, the host interprets the Error Register bits as shown in Table 6-10.

Table 6-10 *Error Register Bits*

MNEMONIC	BIT	DESCRIPTION
BBK	7	Bad block detected in the required sector's ID field.
UNC	6	Uncorrectable data error encountered.
-	5	Not used.
IDNF	4	Requested sector's ID field not found.
-	3	Not used.
ABRT	2	Requested command aborted due to a drive status error, such as Not Ready or Write Fault, or because the command code is invalid.
TK0NF	1	Track 0 not found during execution of RECALIBRATE command.
AMNF	0	Data Address Mark not found after correct ID field format.

6.6.2.3 Sector Count Register

The Sector Count Register defines the number of sectors of data to be transferred across the host bus for a subsequent command. If the value in this register is 0, the sector count is 256 sectors. If the Sector Count Register command executes successfully, the value in this register at command completion is 0. As the drive transfers each sector, it decrements the Sector Count Register to reflect the number of sectors remaining to be transferred. If the command's execution is unsuccessful, this register contains the number of sectors that must be transferred to complete the original request.

When the drive executes an INITIALIZE DRIVE PARAMETERS or Format Track command, the value in this register defines the number of sectors per track.

6.6.2.4 Sector Number Register

The Sector Number Register contains the ID number of the first sector to be accessed by a subsequent command. The sector number is a value between one and the maximum number of sectors per track. As the drive transfers each sector, it increments the Sector Number Register. See the command descriptions in section 6.7 for information about the contents of the Sector Number Register after successful or unsuccessful command completion.

In LBA mode, this register contains bits 0 to 7. At command completion, the host updates this register to reflect the current LBA bits 0 to 7.

6.6.2.5 Cylinder Low Register

The Cylinder Low Register contains the eight low-order bits of the starting cylinder address for any disk access. On multiple sector transfers that cross cylinder boundaries, the host updates this register when command execution is complete, to reflect the current cylinder number. The host loads the least significant bits of the cylinder address into the Cylinder Low Register.

In LBA mode, this register contains bits 8 to 15. At command completion, the host updates this register to reflect the current LBA bits 8 to 15.

6.6.2.6 Cylinder High Register

The Cylinder High Register contains the eight high-order bits of the starting cylinder address for any disk access. On multiple sector transfers that cross cylinder boundaries, the host updates this register at the completion of command execution, to reflect the current cylinder number. The host loads the most significant bits of the cylinder address into the Cylinder High Register.

In LBA mode, this register contains bits 16 to 23. At command completion, the host updates this register to reflect the current LBA bits 16 to 23.

6.6.2.7 Drive/Head Register

The Drive/Head Register contains the drive ID number and its head numbers. By executing the INITIALIZE DRIVE PARAMETERS command, the host defines the contents of the Drive/Head Register.

In LBA mode, this register contains bits 24 to 27. At command completion, the host updates this register to reflect the current LBA bits 24 to 27.

Table 6-11 shows the Drive/Head Register bits.

Table 6-11 *Drive Head Register Bits*

MNEMONIC	BIT	DESCRIPTION
Reserved	7[1]	Always *1*
L	6[2]	*0* for CHS mode *1* for LBA mode
Reserved	5	Always *1*
DRV	4[3]	*0* indicates the Master drive is selected *1* indicates the Slave drive is selected
HS3	3[4]	Most significant Head Address bit in CHS mode Bit 27 of the LBA Address in LBA mode
HS2	2	Head Address bit for CHS mode Bit 26 of the LBA Address in LBA mode
HS1	1	Head Address bit for CHS mode Bit 25 of the LBA Address in LBA mode
HS0	0	Least significant Head Address bit in CHS mode Bit 24 of the LBA Address in LBA mode

1. Bits 5–7 define the sector size set in hardware (512 bytes).
2. Bit 6 is the binary encoded Address Mode Select. When bit 6 is set to 0, addressing is by CHS mode. When bit 6 is set to 1, addressing is by LBA mode.
3. Bit 4 (DRV) contains the binary encoded drive select number. The Master is the primary drive; the Slave is the secondary drive
4. In CHS mode, bits 0–3 (HS0–HS3) contain the binary encoded address of the selected head. At command completion, the host updates these bits to reflect the address of the head currently selected. In LBA mode, bits 0–3 (HS0–HS3) contain bits 24–27 of the LBA Address. At command completion, the host updates this register to reflect the current LBA bits 24 to 27.

6.6.2.8 Status Register

The Status Register contains information about the status of the drive and the controller. The drive updates the contents of this register at the completion of each command. When the Busy bit is set (BSY=1), no other bits in the Command Block Registers are valid. When the Busy bit is not set (BSY=0), the information in the Status Register and Command Block Registers is valid.

When an interrupt is pending, the drive considers that the host has acknowledged the

interrupt when it reads the Status Register. Therefore, whenever the host reads this register, the drive clears any pending interrupt. Table 6-12 defines the Status Register bits.

6.6.2.9 Command Register

The host sends a command to the drive by means of an 8-bit code written to the Command Register. As soon as the drive receives the command in its Command Register, it begins execution of the command. Table 6-13 lists the hexadecimal command codes and parameters for each executable command. The code F0h is common to all of the extended commands. Each of these commands is distinguished by a unique subcode. For a detailed description of each command, see Section 6.7, "COMMAND DESCRIPTIONS," found later in this chapter.

Table 6-12 *Status Register Bits*

MNEMONIC	BIT	DESCRIPTION
BSY	7	Busy bit. Set by the controller logic of the drive whenever the drive has access to and the host is locked out of the Command Block Registers. BSY is set under the following conditions: • Within 400 ns after the deassertion of RESET- or after SRST is set in the Device Control Register. Following a reset, BSY will be set for no longer than 30 seconds. • Within 400 ns of a host write to the Command Block Registers with a Read, READ LONG, READ BUFFER, SEEK, RECALIBRATE, INITIALIZE DRIVE PARAMETERS, Read Verify, Identify Drive, or EXECUTE DRIVE DIAGNOSTIC command. • Within 5 µsec after the transfer of 512 bytes of data during the execution of a Write, Format Track, or WRITE BUFFER command, or 512 bytes of data and the appropriate number of ECC bytes during the execution of a WRITE LONG command. When BSY=1, the host cannot write to a Command Block Register and reading any Command Block Register returns the contents of the Status Register.
DRDY	6	Drive Ready bit. Indicates that the drive is ready to accept a command. When an error occurs, this bit remains unchanged until the host reads the Status Register, then again indicates that the drive is ready. At power on, this bit should be cleared, and should remain cleared until the drive is up to speed and ready to accept a command.
DWF	5	Drive Write Fault bit. When an error occurs, this bit remains unchanged until the host reads the Status Register, then again indicates the current write fault status.
DSC	4	Drive Seek Complete bit. This bit is set when a seek operation is complete and the heads have settled over a track. When an error occurs, this bit remains unchanged until the host reads the Status Register, when it indicates the current seek-complete status.
DRQ	3	Data Request bit. When set, this bit indicates that the drive is ready to transfer a word or byte of data from the host to the data port.
CORR	2	Corrected Data bit. The drive sets this bit when it encounters and corrects a correctable data error. This condition does not terminate a multisector read operation.
IDX	1	Index bit. This bit is set when the drive detects the index mark, once per disk revolution.
ERR	0	Error bit. When set, this bit indicates that the previous command resulted in an error. The other bits in the Status Register and the bits in the Error Register contain additional information about the cause of the error.

Table 6-13 *Quantum Bigfoot CY 2.1/4.3/6.4 AT Command Codes and Parameters*

COMMAND		PARAMETER					
NAME	CODE	SC	SN	CY	DS	HD	FR
RECALIBRATE	1Xh				V		
READ SECTORS, with retry	20h	V	V	V	V	V	
READ SECTORS, no retry	21h	V	V	V	V	V	
READ LONG, with retry	22h	V	V	V	V	V	
READ LONG, no retry	23h	V	V	V	V	V	
WRITE SECTORS, with retry	30h	V	V	V	V	V	
WRITE SECTORS, no retry	31h	V	V	V	V	V	
WRITE LONG, with retry	32h	V	V	V	V	V	
WRITE LONG, no retry	33h	V	V	V	V	V	
READ VERIFY SECTORS, with retry	40h	V	V	V	V	V	
READ VERIFY SECTORS, no retry	41h	V	V	V	V	V	
FORMAT TRACK	50h	V		V	V	V	
SEEK	7Xh		V	V	V	V	
EXECUTE DRIVE DIAGNOSTIC	90h						
INITIALIZE DRIVE PARAMETERS	91h	V			V	V	
READ MULTIPLE	C4h	V	V	V	V	V	
WRITE MULTIPLE	C5h	V	V	V	V	V	
SET MULTIPLE MODE	C6h	V			V		
READ DMA, with retry	C8h	V	V	V	V	V	
READ DMA, no retry	C9h	V	V	V	V	V	
WRITE DMA, with retry	CAh	V	V	V	V	V	
WRITE DMA, no retry	CBh	V	V	V	V	V	
STANDBY IMMEDIATE	E0h					V	
IDLE IMMEDIATE	E1h					V	
STANDBY MODE (AUTO POWER-DOWN)	E2h	V				V	
IDLE MODE (AUTO POWER-DOWN)	E3h	V				V	
READ BUFFER	E4h				V		
CHECK POWER MODE	E5h	V				V	
SLEEP MODE	E6h					V	
WRITE BUFFER	E8h				V		
IDENTIFY DRIVE	ECh				V		
SET FEATURES	EFh				V		V
READ DEFECT LIST–extended cmnd.	F0h	V	V	V	V	V	
READ CONFIGURATION–extended cmnd.	F0h	V	V	V	V	V	
SET CONFIGURATION–extended cmnd.	F0h	V	V	V	V	V	

Note: The following information applies to Table 6-13:

SC = Sector Count Register
SN = Sector Number Register
CY = Cylinder Low and High Registers
DS = Drive Select bit (Bit 4 of Drive/Head Register)
HD = 3 Head Select Bits (Bits 0–3 of Drive Head Register)
FR = Features Register
V = Must contain valid information for this command

6.7 COMMAND DESCRIPTIONS

The Quantum Bigfoot CY 2.1/4.3/6.4 AT hard disk drives support all standard ATA drive commands. The drive decodes, then executes, commands loaded into the Command Block Register. In applications involving two hard drives, both drives receive all commands. However, only the selected drive executes commands–with the exception of the EXECUTE DRIVE DIAGNOSTIC command, as explained below. The procedure for executing a command on the selected drive is as follows:

1. Wait for the drive to indicate that it is no longer busy (BSY=0).
2. Load the required parameters into the Command Block Register.
3. Activate the Interrupt Enable (-IEN) bit.
4. Wait for the drive to set RDY (RDY=1).
5. Write the command code to the Command Register.

Execution of the command begins as soon as the drive loads the Command Block Register. The remainder of this section describes the function of each command. The commands are listed in the same order they appear in Table 6-13.

6.7.1 Recalibrate—1xh

The RECALIBRATE command moves the read/write heads from any location on the disk to cylinder 0. On receiving this command, the drive sets the BSY bit and issues a seek command to cylinder 0. The drive then waits for the seek operation to complete, updates status, negates BSY, and generates an interrupt. If the drive cannot seek to cylinder 0, it posts the message TRACK 0 NOT FOUND.

6.7.2 Read Sectors—20h (with retry), 21h (without retry)

The Read Sectors command reads from 1 to 256 sectors, beginning at the specified sector. As specified in the Command Block Register, a sector count equal to 0 requests 256 sectors. When the drive accepts this command, it sets BSY and begins execution of the command.

6.7.2.1 Read Long—22h (with retry), 23h (without retry)

When the Long bit is set in the command code, a READ LONG command executes, returning the data and the ECC bytes contained in the data field of the requested sector. During a READ LONG operation, the drive does not check the ECC bytes to determine if a data error of any kind has occurred.

6.7.2.2 Multiple Sector Reads

Multiple sector reads set DRQ. After reading each sector, the drive generates an interrupt when the sector buffer is full and the drive is ready for the host to read the data. Once the host empties the sector buffer, the drive immediately clears DRQ and sets BSY.

If an error occurs during a multiple sector read, the read terminates at the sector in which the error occurred. The Command Block Register contains the cylinder, head, and sector numbers of the sector in which the error occurred. The host can then read the Command Block Register to determine what kind of error has occurred, and in which sector. Whether the data error is correctable or uncorrectable, the drive loads the data into the sector buffer.

6.7.3 Write Sector—30h (with retry), 31h (without retry)

The WRITE SECTOR command writes from 1 to 256 sectors, beginning at the specified sector. As specified in the Command Block Register, a sector count equal to 0 requests 256 sectors. When the drive accepts this command, it sets DRQ and waits for the host to fill the sector buffer with the data to be written to the drive. The drive does not generate an interrupt to start the first buffer-fill operation. Once the buffer is full, the drive clears DRQ, sets BSY, and begins execution of the command.

6.7.3.1 Write Long—32h (with retry), 33h (without retry)

When the Long bit is set in the command code, a WRITE LONG command writes the data and the ECC bytes directly from the sector buffer. The drive does not generate the ECC bytes itself.

6.7.3.2 Multiple Sector Writes

The MULTIPLE SECTOR WRITES command sets DRQ. The drive generates an interrupt whenever the sector buffer is ready to be filled. When the host fills the sector buffer, the drive immediately clears DRQ and sets BSY.

If an error occurs during a multiple sector write operation, the write operation terminates at the sector in which the error occurred. The Command Block Register contains the cylinder, head, and sector numbers of the sector in which the error occurred. The host can then read the Command Block Register to determine what kind of error has occurred, and in which sector.

6.7.4 Read Verify Sectors—40h (with retry), 41h (without retry)

The execution of the READ VERIFY SECTORS command is identical to that of the READ SECTORS command. However, the Read Verify command does not cause the drive to set DRQ, the drive transfers no data to the host, and the Long bit is invalid. On receiving the READ VERIFY command, the drive sets BSY. When the drive has verified the requested sectors, it clears BSY and generates an interrupt. On command completion, the Command Block Registers contain the cylinder, head, and sector numbers of the last sector verified.

If an error occurs during a multiple sector verify operation, the read operation terminates at the sector in which the error occurred. The Command Block Registers contain the cylinder, head, and sector numbers in which the error occurred.

6.7.5 Format Track—50h

The host specifies the track addresses by writing to the Cylinder and Head Registers. When the drive accepts a FORMAT TRACK command, it sets the DRQ bit, then waits for the host to fill the sector buffer. When the buffer is full, the drive clears DRQ, sets BSY, and begins command execution. The contents of the sector buffer are ignored and not written to the disk, and will be ignored.

On the Quantum Bigfoot CY 2.1/4.3/6.4 AT hard disk drives, the FORMAT TRACK command writes zeroes to the data fields in the sectors on the specified logical track. The drive writes no headers at these locations. The Sector Count register contains the number of sectors per track.

6.7.6 Seek—7xh

The SEEK command causes the actuator to seek to the track to which the Cylinder and Drive/Head registers point. When the drive receives this command in its Command Block Registers, it performs the following functions:

1. Sets BSY
2. Initiates the seek operation
3. Resets BSY
4. Sets the Drive Seek Complete (DSC) bit in the Status Register

The drive does not wait for the seek to complete before it sends an interrupt. If the BSY bit is *not* set in the Status Register, the drive can accept and queue subsequent commands while performing the seek. If the Cylinder registers contain an illegal cylinder, the drive sets the ERR bit in the Status Register and the IDNF bit in the Error Register.

6.7.7 Execute Drive Diagnostic—90h

The EXECUTE DRIVE DIAGNOSTIC command performs the internal diagnostic tests implemented on the drive. Drive 0 sets BSY within 400 ns of receiving of the command.

If Drive 1 is present:

- Both drives execute diagnostics.
- Drive 0 waits up to six seconds for drive 1 to assert PDIAG–.
- If drive 1 does not assert PDIAG- to indicate a failure, drive 0 appends 80h with its own diagnostic status.
- If the host detects a drive 1 diagnostic failure when reading drive 0 status, it sets the DRV bit, then reads the drive 1 status.

If Drive 1 is not present:

- Drive 0 reports only its own diagnostic results.
- Drive 0 clears BSY and generates an interrupt.

If drive 1 fails diagnostics, drive 0 appends 80h with its own diagnostic status and loads that code into the Error Register. If drive 1 passes its diagnostics or no drive 1 is present, drive 0 appends 00h with its own diagnostic status and loads that code into the Error Register.

The diagnostic code written to the Error Register is a unique 8-bit code. Table 6-14 lists the diagnostic codes.

Table 6-14 *Diagnostic Codes*

DIAGNOSTIC CODE	DESCRIPTION
01h	No Error Detected
02h	Formatter Device Error
03h	Sector Buffer Error
04h	ECC Circuitry Error
05h	Controlling Microprocessor Error
8Xh	Drive 1 Failed

6.7.8 Initialize Drive Parameters—91h

The INITIALIZE DRIVE PARAMETERS command enables the host to set the logical number of heads and the logical number of sectors per track. On receiving the command, the drive sets the BSY bit, saves the parameters, clears BSY, and generates an interrupt.

The only two register values used by this command are the Sector Count Register, which specifies the number of sectors; and the Drive/Head Register, which specifies the number of heads, minus 1. The DRV bit assigns these values to drive 0 or drive 1, as appropriate.

This command does not check the sector count and head values for validity. If these values are invalid, the drive will not report an error until another command causes an illegal access.

6.7.9 Read Multiple—C4h

The execution of the READ MULTIPLE command is identical to that of the Read Sectors command. However, the READ MULTIPLE command:

- Transfers blocks of data to the host without intervening interrupts
- Requires DRQ qualification of the transfer only at the start of the block—*not* at each sector
- Invalidates the Long bit

The SET MULTIPLE MODE command specifies the block count, or the number of sectors to be transferred as a block. This command executes prior to the READ MULTIPLE command. When the host issues a READ MULTIPLE command, the Sector Count Register contains the number of sectors requested—not the number of blocks or the block count.

If this sector count is not evenly divisible by the block count, the drive transfers as many full blocks as possible to the host, followed by a final partial-block transfer. The partial-block transfer is for *n* sectors, where:

n = (sector count) module (block count)

If the drive attempts execution of a READ MULTIPLE command before executing the SET MULTIPLE MODE command or if READ MULTIPLE commands are disabled, an abort command error occurs.

The drive reports disk errors encountered during READ MULTIPLE commands at the beginning of a block or partial-block transfer. However, the drive still sets DRQ and transfers the data—including any corrupted data.

6.7.10 Write Multiple—C5h

The execution of the WRITE MULTIPLE command is identical to that of the Write Sectors command. However, the WRITE MULTIPLE command:

- Causes the controller to set BSY within 400 ns of accepting the command
- Causes the drive to transfer multiple-sector blocks of data to the drive without intervening interrupts
- Requires DRQ qualification of the transfer only at the start of the block, not at each sector
- Invalidates the Long bit

The SET MULTIPLE MODE command specifies the block count, or the number of sectors

to be transferred as a block. This command executes prior to the WRITE MULTIPLE command. When the host issues a WRITE MULTIPLE command, the Sector Count Register contains the number of sectors requested–not the number of blocks or the block count.

If this sector count is not evenly divisible by the block count, the drive transfers as many full blocks as possible, followed by a final partial-block transfer. The partial-block transfer is for n sectors, where:

$$n = \text{(sector count) module (block count)}$$

If the drive attempts to execute a WRITE MULTIPLE command before executing the SET MULTIPLE MODE command or while WRITE MULTIPLE commands are disabled, an Abort Command error occurs.

During the execution of a WRITE MULTIPLE command, the drive reports all disk errors encountered, following an attempted disk write of the block or partial block. When an error occurs, the WRITE MULTIPLE command ends at the sector that contains the error–even if it is in the middle of a block–and does not transfer subsequent blocks. The drive generates interrupts by setting DRQ at the beginning of each block or partial block.

6.7.11 Set Multiple Mode—C6h

The SET MULTIPLE MODE command enables the controller to perform READ MULTIPLE and WRITE MULTIPLE operations, and establishes the block count for these commands.

Prior to issuing a command, the host should load the Sector Count Register with the number of sectors per block. On receiving this command, the drive sets BSY and checks the contents of the Sector Count Register.

If the Sector Count Register contains a valid value and the controller supports block count, the controller loads the values for all subsequent READ MULTIPLE and WRITE MULTIPLE commands and enables execution of these commands. Any unsupported block count in the register causes an Aborted Command error, and disables execution of READ MULTIPLE and WRITE MULTIPLE commands.

If the Sector Count Register contains a zero value when the host issues the command, READ MULTIPLE and WRITE MULTIPLE commands are disabled. Any unsupported block count in the register causes an aborted command error, and disables READ MULTIPLE and WRITE MULTIPLE commands. After the command is executed, the controller clears BSY. At power on, or after a software or hardware reset, the default mode for the READ MULTIPLE and WRITE MULTIPLE commands is disabled.

6.7.12 Read Buffer—E4h

The READ BUFFER command enables the host to read the current contents of the drive's sector buffer. When the host issues this command, the drive sets BSY, sets up the sector buffer for a read operation, sets DRQ, class BSY, and generates an interrupt. The host then reads up to 512 bytes of data from the buffer.

The drive can synchronize READ BUFFER and WRITE BUFFER commands from the host; that is, sequential READ BUFFER and WRITE BUFFER commands can access the same 512 bytes within the buffer.

6.7.13 Write Buffer—E8h

The WRITE BUFFER command allows the host to write the first 512 bytes of the drive's buffer. On receiving this command in its Command Block Register, the drive sets BSY and prepares for a write operation. When ready, the drive sets DRQ, resets BSY, and generates INTRQ, allowing the host to the buffer.

6.7.14 Power Management Commands

The Quantum Bigfoot CY 2.1/4.3/6.4 AT hard disk drives provide numerous management options. Two important options center around a count down counter known as the automatic power down counter or APD. This counter can trigger one of two power saving events depending on which of the two commands was most to write to the buffer.

- **Standby:** Once a standby command is issued, the drive enters the standby mode. Further, each time the APD counter reaches zero in the future, the drive enters the standby mode, the spindle and actuator motors are off and the heads are parked in the landing zone. Receipt of any command that requires media access causes the drive to exit the standby command and service the host request. Each time the drive executes the standby command, the drive will reenter the standby mode when the APD counter reaches zero.
- **Idle:** Once an idle command is issued, each time the APD counter reaches zero, the drive enters the standby mode. In the standby mode, the actuator and spindle motors are off with the heads locked in the landing area. This is the default setting.

Three commands are available which are not dependent upon the APD counter reaching zero:

- **Sleep:** When a sleep command is received, the drive enters the sleep mode. In the sleep mode, the spindle and actuator motors are off and the heads are latched in the landing zone. Receipt of a reset causes the drive to transition from the sleep to the standby mode.
- **Standby Immediate:** When a standby immediate command is received, the drive immediately enters the standby mode.
- **Idle Immediate:** When an idle immediate command is received, after the first decrement of the APD counter, the drive enters the idle mode.

The sleep, standby immediate, and idle immediate commands differ in a significant way from the standby and idle commands. Specifically, sleep, standby immediate, and idle immediate have a one-time effect and must be reissued each time their effect is desired. In contrast, standby and idle operate in conjunction with the APD counter and stay in effect continually, becoming non-effectual only upon issuance of the other of these two commands. Thus, for example, once the standby command is issued just one time, each time the APD counter reaches zero the drive will enter the standby mode.

Note: The user has the ability to determine the value to which the APD counter is set upon completion of any command. This value is set by writing to the Sector Count Register a number between 12 and 255 just prior to issuance of a standby or idle command. Each increment represents a five-second time interval.

6.7.14.1 Standby Immediate Mode – E0h

The Standby Immediate Mode power command immediately puts the drive in the Standby Mode. Power is removed from the spindle motor (the drive's PCB power remains) and the heads are parked.

6.7.14.2 Idle Immediate Mode – E1h

The Idle Immediate Mode power command immediately puts the drive in the Idle Mode.

6.7.14.3 Standby Mode, Automatic Power-Down – E2h

The Standby Mode, Automatic Power-Down (APD) command immediately puts the drive in the Standby Mode. The Sector Count Register is then examined. If the value in this register is not zero, the Auto Power-Down feature is enabled and will take effect once the countdown timer reaches zero. The valid count range is Table 6-16. Each time the drive is accessed, the countdown timer is reset to the value originally set in the Sector Count Register at the time the Standby Mode-Auto Power Down command was issued.

Note: If the value in the Sector Count Register is zero, the Auto Power-Down feature is disabled.

Table 6-15 *Valid Count Range*

SECTOR COUNT	TIME
1 to 12	1 minute
13 to 240	(Value * 5) seconds
241 to 251	(Value – 240) * 30 seconds
252 to 255	((Value * 5) seconds

6.7.14.4 Idle Mode, Automatic Power-Down – E3h

The Idle Mode, Automatic Power-Down command immediately puts the drive into the Idle Mode. The Sector Count Register is then examined. If the value in this register is not zero, the Auto Power-Down feature is enabled and takes effect once the countdown timer reaches zero. The valid count range is listed in Table 6-16. Each time the drive is accessed, the countdown timer is reset to the value originally set in the Sector Count Register at the time the Idle Mode-Auto Power Down command was issued.

Note: If the value in the Sector Count Register is zero, the Auto Power-Down feature is disabled.

6.7.14.5 Check Power Mode – E5h

The CHECK POWER MODE command writes FFh into the Sector Count Register provided that the drive is in the Idle Mode, even if it is in Automatic Power-Down mode. However, if it is in Standby mode, the drive returns a value of 00h in the Sector Count Register.

6.7.14.6 Sleep Mode – E6h

The Quantum Bigfoot drive considers the Sleep Mode to be the equivalent to the Standby Mode, except that a reset is required before issuing a command requiring media access.

6.7.15 Identify Drive—ECh

The IDENTIFY DRIVE command enables the host to receive parameter information from the drive. When the host issues this command, the drive sets BSY, stores the required parameter information in the sector buffer, sets DRQ, and generates an interrupt. The host then reads the information from the sector buffer. The Identify Drive Parameters Table, shown in Table 6-16, defines the parameter words stored in the buffer. All reserved bits should be zeros. A full explanation of the parameter words is listed below:

Default Logical Cylinders: The number of translated cylinders in the default translation mode.

Number of Logical Heads: The number of translated heads in the default translation mode.

Number of Unformatted Bytes Per Track: The number of unformatted bytes per translated track in the default translation mode.

Number of Unformatted Bytes Per Sector: The number of unformatted bytes per sector in the default translation mode.

Number of Logical Sectors Per Track: The number of sectors per track in the default translation mode.

Serial Number: The contents of this field are left aligned and padded with spaces (20h).

Buffer Type: The contents of this field are as follows:

- 0000h = Not specified
- 0001h = A single-ported, single-sector buffer capable of data transfers either to or from the host or to or from the disk
- 0002h = A dual-ported, multiple-sector buffer capable of simultaneous data transfers either to and from the host, or from the host and the disk
- 0003h = A dual-ported, multiple-sector buffer capable of simultaneous data transfers with read caching
- 0004 - FFFFh = Reserved

Firmware Revision: The contents of this field are left-aligned and padded with spaces (20h).

Model Number: The contents of this field are left-aligned and padded with spaces (20h). The low-order byte appears first in a word.

Table 6-16 *Identify Drive Parameters*

WORDS[1]			PARAMETER DESCRIPTION (Statements below are true if the bit is set to 1)
WORD	BIT	BIT VALUE	
0	15	0	Reserved for nonmagnetic drives
	14	0	Format speed tolerance gap required
	13	0	Track offset option available
	12	0	Data strobe offset option available

WORDS[1]			PARAMETER DESCRIPTION (Statements below are true if the bit is set to 1)
WORD	BIT	BIT VALUE	
	11	0	Rotational speed tolerance is > 0.5%
	10	1	Disk transfer rate > 10 Mbit/s
	9	0	Disk transfer rate > 5 Mbit/s, but < 10 Mbit/s
	8	0	Disk transfer rate <= 5 Mbit/s
	7	0	Reserved for removable-cartridge drive
	6	1	Hard disk drive
	5	0	Spindle motor control option implemented
	4	1	Head-switch time > 15 µs
	3	1	Not MFM-encoded
	2	0	Soft-sectored
	1	1	Hard-sectored
	0	0	Reserved
1		2160AT = 4,092 4320AT = 8,960 6480AT = 13,456	Default logical cylinders
2		0	Reserved
3		2160AT = 16 4320AT = 15 6480AT = 15	Default number of logical heads
4			Number of unformatted bytes per track
5		512	Number of unformatted bytes per sector
6		2160AT = 63 4320AT = 63 6480AT = 63	Default number of logical sectors per track
7–9		5154h	Vendor-unique
10–19			Serial number (20 ASCII characters)[2]
20		3	Buffer type
21		A7h	Buffer size in 512-byte increments
22		4	Number of ECC bytes passed on READ/WRITE LONG commands
23–26			Firmware revision (8 ASCII characters)
27–46			Model number (40 ASCII characters)

WORDS[1]			PARAMETER DESCRIPTION (Statements below are true if the bit is set to 1)
WORD	BIT	BIT VALUE	
47	15–8 7–0	80h 08h	Vendor Unique Maximum number of sectors that can be transferred per interrupt is set to 8 for READ MULTIPLE and WRITE MULTIPLE commands.
48		0	Cannot perform double word I/O
49	15–12 11 10 9 8 7–0	0 1 1 1 1 0	Reserved 1 = I/O Ready is supported 1 = I/O Ready can be disabled 1 = LBA supported 1 = DMA supported Vendor Unique
50		0	Reserved
51	15–8 7–0	2 0	PIO data-transfer cycle timing mode Vendor Unique
52	15–8 7–0	2 0	DMA data-transfer cycle timing mode Vendor Unique
53	15–2 1 0	0 1 1	Reserved 1 = The fields in words 64–70 are valid 1 = The fields in words 54–58 are valid
54			Number of current cylinders
55			Number of current heads
56			Number of current sectors per track
57–58			Current capacity in sectors (CHS mode only)
59	15–9 8 7–0	0 1 n^3	Reserved Multiple sector setting is valid Current setting for number of sectors that can be transferred per interrupt on R/W Multiple commands
60–61		2160AT = 4,124,736 4320AT = 8,467,200 6480AT = 12,715,920	Total number of User Addressable Sectors (LBA Mode only)
62	15–8 7–0	4 7	Single-word DMA transfer mode active (Mode 2) Single-word DMA transfer modes supported (Mode 2)
63	15–8 7–0	4 7	Multiword DMA transfer mode active (Mode 1) Multiword DMA transfer modes supported (Mode 1)
64		3	Advanced PIO Mode is supported
65		120	Minimum multiword DMA transfer cycle time (ns) per word

WORDS[1]			PARAMETER DESCRIPTION (Statements below are true if the bit is set to 1)
WORD	BIT	BIT VALUE	
66		333	Manufacturer's recommended multiword DMA cycle time (ns)
67		333	Manufacturer's PIO cycle time (ns) without flow control
68		120	Manufacturer's PIO cycle time (ns) with flow control

1. The format of an ASCII field specifies that, within a word boundary, the low-order byte appears first.
2. The serial number has the following format: 00QTMTCYJJJLSSSSBBB
where:00 = Placeholders
QT = Quantum
M = Place of manufacture
T = Drive type family (fixed at 12)
C = Drive capacity (1 = 2160 MB, 2 = 4320 MB, 3 = 6480 MB)
Y = Last digit of year drive built
JJJ = Julian date
L = Manufacturing production line
SSSS = Sequence of manufacture
BBB = Blanks (placeholders)
3. n is a variable from zero to 8.

6.7.16 Set Features—EFh

The SET FEATURES command is used by the host to establish certain parameters which control execution of the following drive features:

- 02h - Enable write cache feature
- 03h - Set transfer mode based on value in Sector Count Register
- 55h - Disable read look-ahead feature
- 82h - Disable write cache feature
- AAh - Enable read look-ahead feature

At power-on, or after a reset is accomplished by either the hardware or software, the default mode is 4 bytes of ECC, read look-ahead, and write cache enabled.

A host can choose the transfer mechanism by Set Transfer Mode and specifying a value in the Sector Count register. The upper 5 bits define the type of transfer and the low order 3 bits encode the mode value.

PIO Default Transfer Mode, Disable IORDY00000 001
Single Word DMA Mode x00010 nnn
Multiword DMA Mode x00100 nnn

Where "nnn" is a valid mode number for the associated transfer type.

6.7.17 Read Defect List

The READ DEFECT LIST command enables the host to retrieve the drive's defect list. Prior to issuing the Read Defect List command, the host should issue the Read Defect List Length command, which does not transfer any data. Instead, it stores the length in sectors of the defect list in the Sector Count register (1F2), and the Sector Count Number register (1F3), with the Sector Count register containing the LSB of the two-byte value (see Table 6-17). The defect list length is a fixed value for each Quantum product and can be calculated as follows:

length in sectors = ((maximum number of defects) * 8 + 4) + 511)/512

At the completion of the command, the task file registers 1F2 - 1F6 will contain bytes necessary to execute the Read Defect List command, and the host will only need to write the extended command (F0h) to the Command register (1F7) to proceed with the Read Defect List command execution.

Table 6-17 *READ DEFECT LIST LENGTH Command Bytes*

ADDRESS	VALUE	DEFINITION	VALUE (After)
1F2	**0h**	**Defect List Subcode**	**Length in Sectors (LSB)**
1F3	FFh	Password	Length in Sectors (MSB)
1F4	FFh	Password	FFh
1F5	3Fh	Password	3Fh
1F6	AXh (Drive 0)	Drive Select	AXh = Drive 0
	BXh (Drive 1)	–	BXh = Drive 1
1F7	F0h	Extended Command Code	Status Register

In Table 6-17, registers 1F2h through 1F5h must contain the exact values shown. These values function as a key. The drive issues the ILLEGAL COMMAND message if the bytes are not entered correctly.

The AT Read Defect List command is an extended AT command that enables the host to retrieve the drive's defect list. The host begins by writing to address 1F6h to select the drive. Ten the host writes to addresses 1F2h - 1F5h using values indicated in Table 6-18. When the host subsequently writes the extended command code F0h to address 1F7h, the drives sets BSY, retrieves the defect list, sets DRQ, and resets BSY. The host can now read the requested number of sectors (512 bytes) of data. An INTRQ precedes each sector. Bytes 1F2h and 1F3h contain the two-byte number of sectors that the host expects to read, with address 1F2h containing the LSB (see Table 6-18).

The sector count (1F2h - 1F3h) may vary from product to product and if the wrong value is supplied for a specific product, the drive will issue the ILLEGAL COMMAND message. If the host does not know the appropriate sector count for a specific product, it can issue the Read Defect List Length command described in the previous section to set up the task file for the Read Defect List command.

Table 6-18 *READ DEFECT LIST Command Bytes*

ADDRESS	VALUE	DEFINITION
1F2	Length in Sectors (LSB)	Defect List Subcode
1F3	Length in Sectors (MSB)	Password
1F4	FFh	Password
1F5	3Fh	Password
1F6	AXh = Drive 0	Drive Select
	BXh = Drive 1	–
1F7	F0h	Extended Command Code

Note: Registers 1F2h and 1F3h must contain the transfer length that is appropriate for the specific product, and 1F4h and 1F5h must contain the exact values shown. These values function as a key. The drive issues the message ILLEGAL COMMAND if the bytes are not entered correctly.

Pending defects will be excluded from the list, since no alternative sector is being used as their replacement, and since they may be removed from the drive's internal pending list at a later time. Table 6-17 shows the overall format of the defect list, and Table 6-18 shows the format of the individual defect entries.

Table 6-19 *DEFECT LIST Data Format*

BYTE	DESCRIPTION
0	0
1	1Dh
2	8 * (Number of Defects) (MSB)
3	8 * (Number of Defects) (LSB)
4 - 11	Defect Entry #1
12 - 19	Defect Entry #2
	•
	•

Table 6-20 *DEFECT ENTRY Data Format*

BYTE	DESCRIPTION
0	Defect Cylinder (MSB)
1	Defect Cylinder
2	Defect Cylinder (LSB)
3	Defect Head
4	Defect Sector (MSB)
5	Defect Sector
6	Defect Sector
7	Defect Sector (LSB)

Note: Bytes 4 - 7 will be set to FFh for bad track entries

6.7.18 Configuration

In addition to the SET FEATURES command, the Quantum Bigfoot CY 2.1/4.3/6.4 AT hard disk drives provide two configuration commands:

- The SET CONFIGURATION command, which enables the host to change DisCache and Error Recovery parameters
- The READ CONFIGURATION command, which enables the host to read the current configuration status of the drive

See Chapter 5 for more details about DisCache and setting cache parameters. See Chapter 5 also for more information about error detection and defect management.

6.7.18.1 Read Configuration

The READ CONFIGURATION command displays the configuration of the drive. Like the SET CONFIGURATION command, this command is secured to prevent accidentally accessing it. To access the READ CONFIGURATION command, you must write the pattern shown inTable 6-21 to the Command Block Registers. The first byte, 01h, is a subcode to the extended command code, F0h.

Table 6-21 *Accessing the READ CONFIGURATION Command*

ADDRESS	VALUE	DEFINITION
1F2h	01h	Read Configuration Subcode
1F3h	FFh	Password
1F4h	FFh	Password
1F5h	3Fh	Password
1F6h	AXh (Drive 0)	Drive Select
	BXh (Drive 1)	Drive Select
1F7h	F0h	Extended Command Code

Note: In Table 6-21:

Only the value in address 1F2h of the Command Block Registers is different from the SET CONFIGURATION command.

Registers 1F2h through 1F5h must contain the exact values shown in Table 6-21. These values function as a key. The drive issues the message ILLEGAL COMMAND if the key is not entered correctly.

To select the drive for which the configuration is to be read, set register 1F6h. For execution of the command to begin, load register 1F7h with F0h.

A 512-byte data field is associated with the READ CONFIGURATION command. A 512-byte read sequence sends this data from the drive to the host. The information in this data field represents the current settings of the configuration parameters. The format of the READ CONFIGURATION command data field is similar to that for the data field of the SET CONFIGURATION command, shown in Table 6-22. However, in the READ CONFIGURATION command, bytes 0 through 31 of the data field are *not* KEY information, as they are in the SET CONFIGURATION command. The drive reads these bytes as *QUANTUM CONFIGURATION*, followed by eleven spaces. Users can read the configuration into a buffer, then alter the configuration parameter settings.

6.7.18.2 Set Configuration

The SET CONFIGURATION command is secured to prevent accessing it accidentally. To access the SET CONFIGURATION command, you must write the pattern shown in Table 6-22 to the Command Block Registers. The first byte, FFh, is a subcode to the extended command code F0h.

Table 6-22 *Accessing the SET CONFIGURATION Command*

ADDRESS	VALUE	DEFINITION
1F2h	FFh	Set Configuration Subcode
1F3h	FFh	Password
1F4h	FFh	Password
1F5h	3Fh	Password
1F6h	AXh (Drive 0)	Drive Select
	BXh (Drive 1)	Drive Select
1F7h	F0h	Extended Command Code

Note: In Table 6-22:
Registers 1F2h through 1F5h must contain the exact values shown above. These values function as a key. The drive issues the message ILLEGAL COMMAND if the key is not entered correctly.
To select the drive being reconfigured, register 1F6h should be set.
For execution of the command to begin, load register 1F7h with F0h.

6.7.18.3 Set Configuration Without Saving to Disk

The SET CONFIGURATION WITHOUT SAVING TO DISK command is secured to prevent accidentally accessing it. To access this command, you must write the pattern shown in Table 6-23 to the Command Block Registers. The first byte, FFh, is a subcode to the extended command code F0h.

Table 6-23 *Accessing the SET CONFIGURATION WITHOUT SAVING TO DISK Command*

ADDRESS	VALUE	DEFINITION
1F2h	FEh	Set Configuration Subcode
1F3h	FEh	Password
1F4h	FEh	Password
1F5h	3Fh	Password
1F6h	AXh (Drive 0)	Drive Select
	BXh (Drive 1)	Drive Select
1F7h	F0h	Extended Command Code

Note: In Table 6-23:
Registers 1F2h through 1F5h must contain the exact values shown above. These values function as a key. The drive issues the message ILLEGAL COMMAND if the key is not entered correctly.
To select the drive being reconfigured, set register 1F6h. For execution of the command to begin, load register 1F7h with F0h.

6.7.18.4 Configuration Command Data Field

A 512-byte data field is associated with this command. This data field is sent to the drive through a normal 512-byte write handshake. Table 6-24 shows the format of the data field. Bytes 0 through 31 of the data field contain additional KEY information. The drive issues the message ILLEGAL COMMAND if this information is not entered correctly. Bytes 32 through 35 control the operation of DisCache. Bytes 36 through 38 control operation of the error recovery procedure. The drive does not use bytes 40 through 511, which should be set to 0.

Table 6-24 *Configuration Command Format*

BYTE	BIT							
	7	6	5	4	3	2	1	0
0–31	QUANTUM CONFIGURATION KEY							
32	RESERVED = 0						PE	CE
33	RESERVED = 0							
34	RESERVED = 0							
35	RESERVED = 0							
36	AWRE	ARRE	N/A	RC	EEC	N/A	N/A	DCR
37	NUMBER OF RETRIES							
38	ECC CORRECTION SPAN							
39	RESERVED = 0					WCE	RUEE	N/A
40–511	RESERVED = 0							

Note: All fields marked RESERVED or N/A should be set to zero.

6.7.18.5 Quantum Configuration Key (Bytes 0-31)

Bytes 0-6 must contain the ASCII characters *Q, U, A, N, T, U,* and *M*; byte 7, the ASCII character *space*; and bytes 8-20 must contain the ASCII characters *C, O, N, F, I, G, U, R, A, T, I, O,* and *N*. Bytes 21-31 must contain an ASCII *space*. If this information is not entered correctly, the drive aborts the COMMAND.

6.7.18.6 DisCache Parameters

PE – Prefetch Enable (Byte 32, Bit 1): When set to 1, this bit indicates that the drive will perform prefetching. A PE bit set to 0 indicates that no prefetching will occur. The CE bit (bit 0) must be set to 1 to enable use of the PE bit. The default value is 1.

CE – Cache Enable (Byte 32, Bit 0): When set to 1, this bit indicates that the drive will activate caching on all READ commands. With the CE bit set to 0, the drive will disable caching and use the RAM only as a transfer buffer. The default setting is 1.

6.7.18.7 Error Recovery Parameters

AWRE – Automatic Write Reallocation Enabled (Byte 36, Bit 7): When set to 1, indicates that the drive will enable automatic reallocation of bad blocks. Automatic Write Reallocation is similar to the function of Automatic Read Reallocation, but is initiated by the drive when a defective block has become inaccessible for writing. An AWRE bit set to 0 indicates that the Quantum Bigfoot CY 2.1/4.3/6.4 AT drives will not automatically reallocate bad blocks. The default setting is 1.

ARRE – Automatic Read Reallocation Enable (Byte 36, Bit 6): When set to 1, this bit indicates that the drive will enable automatic reallocation of bad sectors. The drive initiates reallocation when the ARR bit is set to 1 and the drive encounters a hard error–that is, if the triple-burst ECC algorithm is invoked. The default setting is 1. When the ARR bit is set to 0, the drive will not perform automatic reallocation of bad sectors. If RC (byte 36, bit 4) is 1, the drive ignores this bit. The default value is 1.

RC – Read Continuous (Byte 36, Bit 4): When set to 1, this bit instructs the drive to transfer data of the requested length without adding delays to increase data integrity–that is, delays caused by the drive's error-recovery procedures. With RC set to 1 to maintain a continuous flow of data and avoid delays, the drive may send data that is erroneous. When the drive ignores an error, it does *not* post the error. The RC bit set to 0 indicates that potentially time-consuming operations for error recovery are acceptable during data transfer. The default setting is 0.

EEC – Enable Early Correction (Byte 36, Bit 3): When set to 1, this bit indicates that the drive will use its ECC algorithm if it detects two consecutive equal, nonzero error syndromes. The drive will not perform rereads before applying correction, unless it determines that the error is uncorrectable. An EEC bit set to 0 indicates that the drive will use its normal recovery procedure when an error occurs: rereads, followed by error correction. If the RC bit (byte 36, bit 4) is set to 1, the drive ignores the EEC bit. The default setting is 0.

DCR – Disable Correction (Byte 36, Bit 0): When set to 1, this bit indicates that all data will be transferred without correction, even if it would be possible to correct the data. A DCR bit set to 0 indicates that the data will be corrected if possible. If the data is uncorrectable, it will be transferred without correction, though the drive will attempt rereads. If RC (byte 36, bit 4) is set to 1, the drive ignores this bit. The default setting is 0. The drive will post all errors, whether DCR is set to 0 or 1.

NUMBER OF RETRIES (Byte 37): This byte specifies the number of times that the drive will attempt to recover from data errors by rereading the data, before it will apply correction. The drive performs rereads before ECC correction–unless EEC (byte 36, bit 3) is set to 1, enabling early correction. The default is eight.

ECC CORRECTION SPAN (Byte 38): This byte specifies the maximum number of bits per interleave that can be corrected using triple-burst ECC. The value for this byte is fixed at 96.

6.7.18.8 Drive Parameters

WCE – Write Cache Enable (Byte 39, Bit 2): When this bit is set to 1, the Quantum Bigfoot CY 2.1/4.3/6.4 AT hard disk enables the Write Cache. This indicates that the drive returns GOOD status for a write command after successfully receiving the data, but before writing it to the disk. A value of zero indicates that the drive returns GOOD status for a write command after successfully receiving the data and writing it to the disk.

If the next command is another WRITE command, cached data continues to be written to the disk while new data is added to the buffer. The default setting is 1.

RUEE – Reallocate Uncorrectable Error Enables (Byte 39, Bit 1): When set to 1, this bit indicates that the Quantum Bigfoot CY 2.1/4.3/6.4 AT hard disk drive will automatically reallocate uncorrectable hard errors, if the ARR bit (byte 36, bit 6) is set to 1. The default setting is 1.

6.8 ERROR REPORTING

At the start of a command's execution, the Quantum Bigfoot CY 2.1/4.3/6.4 AT hard disk drives check the Command Register for any conditions that would lead to an abort command error. The drive then attempts execution of the command. Any new error causes execution of the command to terminate at the point at which it occurred. Table 6-25 lists the valid errors for each command.

Table 6-25 *Command Errors*

COMMAND	ERROR REGISTER						STATUS REGISTER				
	BBK	UNC	IDNF	ABRT	TKO	AMNF	DRDY	DWF	DSC	CORR	ERR
Check Power Mode				V			V	V	V		V
Read Defect List	V	V	V	V		V	V	V	V	V	V
Execute Drive Diag.											V
Format Track			V	V			V	V	V		V
Identify Drive				V			V	V	V		V
Initialize Parameters							V	V	V		
Invalid Cmnd. Codes				V			V	V	V		V
Read Buffer				V			V	V	V		V
Read DMA	V	V	V	V		V	V	V	V	V	V
Read Configuration	V	V	V	V		V	V	V	V	V	V
Read Multiple	V	V	V	V		V	V	V	V	V	V
Read Sectors	V	V	V	V		V	V	V	V	V	V
Read Sectors Long	V		V	V		V	V	V	V		V
Read Verify Sectors	V	V	V	V		V	V	V	V	V	V
Recalibrate				V	V		V	V	V		V
Seek			V	V			V	V	V		V
Set Configuration	V		V	V			V	V	V		V
Set Features				V			V	V	V		V
Set Multiple Mode				V			V	V	V		V
Write Buffer				V			V	V	V		V
Write DMA	V		V	V			V	V	V		V
Write Multiple	V		V	V			V	V	V		V
Write Sectors	V		V	V			V	V	V		V
Write Sectors Long	V		V	V			V	V	V		V

Note: V = Valid errors for each command
ABRT = Abort command error
AMNF= Data address mark not found error
BBK = Bad block detected
CORR= Corrected data error
DRDY= Drive ready
DSC = Disk seek complete not detected
DWF = Drive write fault detected
ERR = Error bit in the Status Register
IDNF = Requested ID not found
TK0 = Track zero not found error
UNC = Uncorrectable data error

Glossary

A

ACCESS – (v) Read, write, or update information stored on a disk or other medium. (n) The operation of reading, writing, or updating stored information.

ACCESS TIME – The interval between the time a request is made by the system and the time the data is available from the drive. Includes the seek time, rotational latency, and command processing overhead time. (See also *seek*, *rotational latency*, and *overhead*.)

ACTUATOR – Also known as the *positioner*. The internal mechanism that moves the read/write head to the proper track. The Quantum actuator consists of a rotor connected to head mounting arms that position the heads over the desired cylinder. Also known as rotary actuator.

AIRLOCK – A patented Quantum feature that ensures durable and reliable data storage. Upon removal of power from the drive for any reason, the read/write heads automatically park and lock in a non data area called the landing zone. AIRLOCK allows the drive to withstand high levels of non-operating shock. When power is applied to the drive, airflow created from the spinning disks causes the AIRLOCK arm to swing back and unlock the actuator, allowing the heads to move from the landing zone. Upon power down, the AIRLOCK swings back to the locked position, locking the heads in the landing zone. A park utility is not required to park the heads on drives equipped with AIRLOCK (all Quantum drives).

ALLOCATION – The process of assigning particular areas of the disk to specific data or instructions. An allocation unit is a group of sectors on the disk reserved for specified information. On hard disks for small computer systems, the allocation unit is usually in the form of a sector, block, or cluster. (See also *allocation unit*.)

ALLOCATION UNIT – An allocation unit, also known as a *cluster*, is a group of sectors on the disk that can be reserved for the use of a particular file.

ASIC – Acronym for *Application Specific Integrated Circuit*.

AT – An interface designed for IBM PCs and compatible systems. Also known as IDE. (See also *interface*.)

AVERAGE SEEK TIME – The average time it takes for the read/write head to move to a specific location. Calculated by dividing the time it takes to complete a large number of random seeks by the number of seeks performed.

B

BACKUP – A copy of a file, directory, or volume on a separate storage device from the original, for the purpose of retrieval in case the original is accidentally erased, damaged, or destroyed.

BAD BLOCK – A block (usually the size of a sector) that cannot reliably hold data due to a physical flaw or damaged format markings.

BAD TRACK TABLE – A label affixed to the casing of a hard disk drive stating which tracks are flawed and cannot hold data. This list is typed into the low-level formatting program when the drive is installed. Quantum users can ignore bad track tables since Quantum's built-in defect-management protections compensate for these flaws automatically.

BEZEL – A plastic panel that extends the face of a drive so that it covers a computer's drive bay opening. The bezel usually contains a drive-activity LED. Also known as the *faceplate*.

BIOS – Acronym for Basic Input / Output System. The firmware portion of a computer that manages the flow of signals through the system bus and to the attached cards and peripheral devices.

BIT – Abbreviation for binary digit. A binary digit may have one of two values–1 or 0. This contrasts with a decimal digit, which may have a value from 0 to 9. A bit is one of the logic 1or logic 0 binary settings that make up a byte of data. (See also *byte*.)

BLOCK – In UNIX workstation environments, the smallest contiguous area that can be allocated for the storage of data. UNIX blocks are generally 8 Kbytes (16 sectors) in size. In DOS environments, the block is referred to as a cluster. (Note: This usage of the term block at the operating system level is different from its meaning in relation to the physical configuration of the hard drive. See *sector* for comparison.) While an operating system may have its block size set at 8K, the drive usually assigns its "logical block" size to 1 sector (1 OS block = 16 drive logical blocks).

BPI – Bits Per Inch. A measure of how densely information is packed on a storage medium. (See also *FCI*.)

BUFFER – An area of RAM reserved for temporary storage of data that is waiting to be sent to a device that is not yet ready to receive it. The data is usually on its way to or from the hard disk drive or some other peripheral device.

BUS – The part of a chip, circuit board, or interface designed to send and receive data.

BYTE – The basic unit of computer memory, large enough to hold one character of alphanumeric data. Comprised of eight bits. (See also *bit*.)

C

CACHE – High-speed RAM used as a buffer between the CPU and the hard disk. Since the CPU can get information more quickly from the cache than from main memory, the cache usually contains information that is used frequently by the system.

CAPACITY – The amount of information that can be stored on a hard drive. Also known as storage capacity. (See also *formatted capacity*.)

CLEAN ROOM – An environmentally controlled dust-free assembly or repair facility in which hard disk drives are assembled or can be opened for internal servicing.

CLUSTER – In DOS environments, the smallest contiguous area that can be allocated for the storage of data. DOS clusters are usually 2 Kbytes (4 sectors) in size.

CONTROLLER – The chip or circuit that translates computer data and commands into a form suitable for use by the hard drive. Also known as disk controller.

CONTROLLER CARD – An adapter containing the control electronics for one or more hard disks. Usually installed in a slot in the computer.

CPU – Central Processing Unit. The microprocessor chip that performs the bulk of data processing in a computer.

CRC – Cyclic Redundancy Check. An error detection procedure that identifies incomplete or faulty data in each sector.

CYLINDER – When disks are placed directly above one another along the shaft, the circular, vertical "slice" consisting of all the tracks located in a particular position.

D

DATA SEPARATOR – The circuit that extracts data from timing information on drives that store a combined data and clock signal.

DEDICATED SERVO – A positioning mechanism using a dedicated surface of the disk that contains timing and positioning information only, as compared to surfaces that are also used for data. (See also *embedded servo*.)

DEFECT MANAGEMENT – A technique ensuring long-term data integrity. Consists of scanning disk drives both at the factory and during regular use, de-allocating defective sectors before purchase and compensating for new defective sectors afterward.

DISK – In general, any circular-shaped data-storage medium that stores data on the flat surface of the platter. The most common type of disk is the magnetic disk, which stores data as magnetic patterns in a metal or metal-oxide coating. Magnetic disks come in two forms: floppy and hard. Optical recording is a newer disk technology that gives higher capacity storage but at slower access times.

DISK CONTROLLER – A plug-in board, or embedded circuitry on the drive, that passes information to and from the disk. The Quantum hard disk drives all have controllers embedded on the drive printed-circuit board. (See also *controller*.)

DMA – Direct Memory Access. A process for transferring data directly to and from main memory, without passing through the CPU. DMA improves the speed and efficiency by allowing the system to continue processing even while new data is being retrieved.

DOS – Disk Operating System. The most common operating system used in IBM PCs. Manages all access to data on the disk.

DRIVE – Short form of *disk drive*.

DRIVE GEOMETRY – The functional dimensions of a drive, including the number of heads, cylinders, and sectors per track. (See also *logical format*.)

E

ECC – Error Correction Code. The incorporation of extra parity bits in transmitted data in order to detect errors that can be corrected by the controller.

EISA – Extended Industry Standard Architecture. An enhanced AT bus architecture designed by nine manufacturers of PC compatibles and announced in September 1988. EISA promises the advantages of IBM Micro Channel Architecture with the advantage of downward compatibility to the current standard AT bus. (See also *ISA*.)

EMBEDDED SERVO – A timing or location signal placed on tracks that store data. These signals allow the actuator to fine-tune the position of the read/write heads.

ENCODING – The conversion of data into a pattern of On/Off or 1/0 signals prior to being written on the disk surface. (See also *RLL* and *MFM*.)

EPROM – Erasable Programmable Read-Only Memory. An integrated circuit memory chip that can store programs and data in a non-volatile state. These devices can be erased by ultraviolet light and reprogrammed with new data.

EXTERNAL DRIVE – A drive mounted in an enclosure separate from the computer system enclosure, with its own power supply and fan, and connected to the system by a cable.

F

FAT – File Allocation Table. A data table stored on the outer edge of a disk, telling the operating system which sectors are allocated to each file and in what order.

FCI – Flux Changes per Inch. The number of magnetic field patterns that can be stored on a given area of disk surface, used as a measure of data density. (See also *BPI*.)

FILE SERVER – A computer that provides network stations with controlled access to shareable resources. The network operating system is loaded on the file server, and most shareable devices (disk subsystems, printers) are attached to it. The file server controls system security and monitors station-to-station communications. A dedicated file server can be used only as a file server while it is on the network. A non dedicated file server can be used simultaneously as a file server and a workstation.

FIRMWARE – Permanent instructions and data programmed directly into the circuitry of read-only memory for controlling the operation of the computer. Distinct from software, which can be altered by programmers.

FLUX DENSITY – The number of magnetic field patterns that can be stored in a given length of disk surface. The number is usually stated as flux changes per inch (FCI), with typical values in the thousands. (See also *FCI*.)

FLYING HEIGHT – The distance between the read/write head and the disk surface, made up of a cushion of air that keeps the two objects from touching. Smaller flying heights permit denser data storage but require more precise mechanical designs. Also known as fly height.

FORMAT – To write a magnetic track pattern onto a disk surface, specifying the locations of the tracks and sectors. This information must exist on a disk before it can store data.

FORMATTED CAPACITY – The amount of room left to store data on a disk after writing sector headers, boundary definitions, and timing information during a format operation. The size of a Quantum drive is always expressed in formatted capacity, accurately reflecting the usable space required.

FORM FACTOR – The industry standard that defines the physical, external dimensions of a particular device. For example, most Quantum hard disk drives use a 3 1/2-inch form factor.

G

GUIDE RAILS – Plastic strips attached to the sides of a hard disk drive in an IBM PC/AT or compatible computer so that the drive easily slides into place.

H

HALF-HEIGHT – Standard drive size equivalent to half the vertical space of a 5 1/4-inch drive.

HARD DISK – A type of storage medium that retains data as magnetic patterns on a rigid disk, usually made of an iron oxide or alloy over a magnesium or aluminum platter. Because hard disks spin more rapidly than floppy disks, and the head flies closer to the disk, hard disks can transfer data faster and store more in the same volume.

HARD ERROR – A data error that persists when the disk is re-read, usually caused by defects in the physical surface.

HARD-SECTORED – The most common method of indicating the start of each sector on a disk, based on information located in the embedded servo. This method is more precise than soft-sectored techniques and results in lower overhead. (See also *soft-sectored.*)

HEAD – The tiny electromagnetic coil and metal pole used to create and read back magnetic patterns on the disk. Also known as read/write head.

HEAD CRASH – Damage to the read/write head, usually caused by sudden contact with the disk surface. Head crash can also be caused by dust and other particles.

HIGH-CAPACITY DRIVE – By industry conventions typically a drive of 100 megabytes or more.

HIGH-LEVEL FORMATTING – Formatting performed by the operating system to create the root directory, file allocation tables and other basic configurations. (See also *low-level formatting.*)

HOME – Reference track used for recalibration of the actuator. Usually the outermost track (track 0).

HOST ADAPTER – A plug-in board that acts as the interface between a computer system bus and the disk drive.

I

IDE – Integrated Device Electronics. The term used for the interface on drives with embedded controllers that can be directly connected to the AT system bus. This connection can be through a 40-pin connector on the motherboard or an adapter board that contains only the drive-select decoding logic. Quantum AT drives use the IDE interface.

INITIALIZATION – See *low-level formatting.*

INTERFACE – A hardware or software protocol, (contained in the electronics of the disk controller and disk drive) that manages the exchange of data between the drive and computer. The most common interfaces for small computer systems are AT (also known as IDE) and SCSI.

INTERLEAVE – The arrangement of sectors on a track. The Interleave Factor is the number of sectors that pass beneath the read/write heads before the next sector arrives. For example, a 3:1 interleave factor means that the heads read a sector, then let two pass by before reading another, requiring three full revolutions of the disk to access the complete data track. Quantum drives have an interleave factor of 1:1, allowing the system to access a full track of data in a single revolution.

INTERLEAVE FACTOR – The number of sectors that pass beneath the read/write heads before the next numbered sector arrives. When the interleave factor is 3:1, a sector is read, two pass by, and then the next is read. It would take three revolutions of the disk to access a full track of data. Quantum drives have an interleave of 1:1, so a full track of data can be accessed within one revolution of the disk, thus offering the highest data throughput possible.

INTERNAL DRIVE – A drive mounted inside one of a computer's drive bays, or a hard disk on a card installed in one of the computer's expansion slots.

ISA – Industry Standard Architecture. The standard 16 bit AT bus as designed by IBM for their PC/AT system. This has been the only industry standard bus for PCs until the recent release of MCA (Micro Channel Architecture) and EISA (Extended Industry Standard Architecture). (See also *EISA.*)

J

JUMPER – A tiny box that slips over two pins on a circuit board, connecting the pins electrically. Some board manufacturers use Dual In-Line Package (DIP) switches instead of jumpers.

K

KILOBYTE (K) – A unit of measure consisting of 1,024 (2^{10}) bytes.

L

LANDING ZONE – A non-data area on the disk's inner cylinder where the heads can rest when the power is off.

LATENCY – The time during which the read/write heads wait for the data to rotate into position after the controller starts looking for a particular data track. If a disk rotates at 3,600 rpm, the maximum latency time is 16.4 milliseconds, and the average latency time is 8.2 milliseconds.

LOGICAL BLOCK – See Block.

LOGICAL FORMAT – In connection with standard disk formatting, refers to low-level formatting. In relation to DOS-specific format requirements, refers to the translations accomplished by the controller in situations where the hard drive data configurations do not match DOS format limitations.

LOOK AHEAD – The process of anticipating events in order to speed up computer operations. For example, the system can buffer data into cache RAM by reading blocks in advance, preparing the system for the next data request.

LOW-LEVEL FORMATTING – The process of creating sectors on the disk surface so that the operating system can access the required areas for generating the file structure. Quantum drives are shipped with the low-level formatting already completed. Also known as *initialization*.

M

MB – See *megabyte*.

MEDIA – The magnetic film that is deposited or coated on an aluminum substrate which is very flat and in the shape of a disk. The media is overcoated with a lubricant to prevent damage to the heads or media during head take off and landing. The media is where the data is stored inside the disk in the form of magnetic flux or polarity changes.

MEGABYTE (MB) –Quantum defines megabyte as 10^6 bytes or 1,000,000 bytes.

MEGAHERTZ – A measurement of frequency in millions of cycles per second.

MHz – See *megahertz*.

MICROPROCESSOR – The integrated circuit chip that performs the bulk of data processing and controls the operation of all of the parts of the system. A disk drive also contains a microprocessor to handle all of the internal functions of the drive and to support the embedded controller.

MICROSECOND (µs) – One millionth of a second (.000001 sec.).

MILLISECOND (ms) – One thousandth of a second (.001 sec.).

MR HEAD – An MR (Magnetoresistive) head employs two closely-spaced but separate write and read elements: a thin-film inductive element for recording data, and a magnetically sensitive thin-film resistive element for reading data. Each of the two elements can be optimized to perform its particular function, reading or writing data.

MTBF – Mean Time Between Failure. Reliability rating indicating the failure rate expected of a product expressed in power on hours (POH). Since manufacturers differ in the ways they determine the MTBF, comparisons of products should always take into account the MTBF calculation method.

MTTR – Mean Time To Repair. The average time it takes to repair a drive that has failed for some reason. This only takes into consideration the changing of the major sub-assemblies such as circuit board or sealed housing. Component level repair is not included in this number as this type of repair is not performed in the field.

O

OVERHEAD – *Command overhead* refers to the processing time required by the controller, host adapter, or drive prior to the execution of a command. Lower command overhead yields higher drive performance. (See also *zero command overhead.*) *Disk overhead* refers to the space required for non-data information such as location and timing. Disk overhead often accounts for about ten percent of drive capacity. Lower disk overhead yields greater disk capacity.

OVERWRITE – To write data on top of existing data, erasing it.

OXIDE – A metal-oxygen compound. Most magnetic coatings are combinations of iron or other metal oxides, and the term has become a general one for the magnetic coating on tape or disk.

P

PARTITION – A portion of a hard disk dedicated to a particular operating system and application and accessed as a single logical volume.

PERFORMANCE – A measure of the speed of the drive during normal operation. Factors affecting performance are seek times, transfer rate and command overhead.

PERIPHERAL – A device added to a system as an enhancement to the basic CPU, such as a disk drive, tape drive or printer.

PHYSICAL FORMAT – The actual physical layout of cylinders, tracks, and sectors on a disk drive.

PLATED MEDIA – Disks that are covered with a hard metal alloy instead of an iron-oxide compound. Plated disks can store greater amounts of data than their oxide-coated counterparts.

PLATTER – Common term referring to the hard disk.

POH – Power On Hours. The unit of measurement for Mean Time Between Failure as expressed in the number of hours that power is applied to the device regardless of the amount of actual data transfer usage. (See also *MTBF*.)

POSITIONER – See *actuator*.

PROGRAMMED I/O – In a disk drive with an AT interface, data may be transferred between the drive and host using programmed I/O (PIO). In this case, the registers in the disk drive are accessed using host address lines A0-A2, chip select signals CS1FX-and CS3FX-, and read/write signals IOR- a nd IOW-. The host uses PIO to write to the Command Block Registers when transmitting commands to the drive, and to the Control Block Registers when transmitting control, such as a software reset.

R

RAM – Random Access Memory. An integrated circuit memory chip that allows information to be stored and retrieved by a microprocessor or controller. The information may be stored and retrieved in any order, and all storage locations are equally accessible.

RAM DISK – A "phantom" disk drive created by setting aside a section of RAM as if it were a group of regular sectors. Access to RAM disk data is extremely fast, but is lost when the system is reset or turned off.

READ AFTER WRITE – A mode of operation requiring that the system read each sector after data is written, checking that the data read back is the same as the data recorded. This operation lowers system speed but raises data reliability.

READ VERIFY – A data accuracy check performed by having the disk read data to the controller, which then checks for errors but does not pass the data on to the system.

READ/WRITE HEAD – The tiny electromagnetic coil and metal pole piece used to create and read back the magnetic patterns (write or read information) on the disk. Each side of each platter has its own read/write head.

REMOVABLE DISK – Generally said of disk drives where the disk itself is meant to be removed, and in particular of hard disks using disks mounted in cartridges. Their advantage is that multiple disks can be used to increase the amount of stored material, and that once removed, the disk can be stored away to prevent unauthorized use.

RLL – Run Length Limited. A method of encoding data into magnetic pulses. The RLL technique permits 50% more data per disk than the MFM method, but requires additional processing.

ROM – Read-Only Memory. Integrated circuit memory chip containing programs that can be accessed and read but can not be modified.

ROTARY ACTUATOR – The rotary actuator replaces the stepper motor used in the past by many hard disk manufacturers. The rotary actuator is perfectly balanced and rotates around a single pivot point. It allows closed-loop feedback positioning of the heads, which is more accurate than stepper motors.

ROTATIONAL LATENCY – The delay between when the controller starts looking for a specific block of data on a track and when that block rotates around to where it can be read by the read/write head. On average, it is half of the time needed for a full rotation (about 8 ms.).

S

SECTOR – On a PC hard drive, the minimum segment of track length that can be assigned to store information. On Macintosh and UNIX drives, sectors are usually grouped into blocks or logical blocks that function as the smallest data unit permitted. Since these blocks are often defined as a single sector the terms block and sector are sometimes used interchangeably in this

context. (Note: The usage of the term block in connection with the physical configuration of the disk is different from its meaning at the system level. See also *block* and *cluster* for comparison.)

SEEK - A movement of the disk read/write head to a specific data track.

SERVO DATA - Magnetic markings written on the media that guide the read/write heads to the proper position.

SERVO SURFACE - A separate surface containing only positioning and disk timing information but no data.

SETTLE TIME - The interval between the arrival of the read/write head at a specific track, and the lessening of the residual movement to a level sufficient for reliable reading or writing.

SHOCK RATING - A rating, expressed in "G's", of how much shock a disk drive can sustain without damage.

SOFT ERROR - A faulty data reading that does not recur if the same data is reread from the disk, or corrected by ECC. Usually caused by power fluctuations or noise spikes.

SOFT-SECTORED - Old time-based method of indicating the start of each sector on a disk. Soft-sectored drives require that location instructions be located in the data fields. (See also *hard-sectored.*)

SPINDLE - The drive's center shaft, on which the hard disks are mounted. A synchronized spindle is a shaft that allows two disks to spin simultaneously as a mirror image of each other, permitting redundant storage of data.

SPUTTER - A special method of coating the disk that results in a hard, smooth surface capable of storing data at a high density. Quantum disk drives use sputtered thin film disks.

STEPPER - A type of motor that moves in discrete with each electrical pulse. Stepper were originally the most common type of actuator engine, since they can be geared to advance a read/write head one track per step. However, they are not as fast, reliable, or durable as the voice coil actuators found in Quantum disk drives. (See also *voice coil.*)

SUBSTRATE - The material underneath the magnetic coating of a disk. Common substrates include aluminum or magnesium alloys for hard drives, glass, for optical disks, and mylar for floppy disks.

SURFACE - The top or bottom side of a disk, which is coated with the magnetic material for recording data. On some drives one surface may be reserved for positioning information.

T

THIN FILM - A type of coating allowing very thin layers of magnetic material, used on hard disks and read/write heads. Hard disks with thin film surfaces can store greater amounts of data.

TPI - Tracks Per Inch. The number of tracks written within each inch of disk's surface, used as a measure of how closely the tracks are packed on a disk surface. Also known as *track density.*

TRACK - One of the many concentric magnetic circle patterns written on a disk surface as a guide for storing and reading data. Also known as *channel.*

TRACK DENSITY - How closely the tracks are packed on a disk surface. The number is specified as tracks per inch (TPI).

TRACK-TO-TRACK SEEK TIME - The time required for the read/write heads to move to an adjacent track.

TRANSFER RATE - The rate at which the disk sends and receives data from the controller. The sustained transfer rate includes the time required for system processing, head switches and seeks, and accurately reflects the drive's true performance. The burst mode transfer rate is a much higher figure that refers only to the movement of data directly into RAM.

U

UNFORMATTED CAPACITY - The total number of usable bytes on a disk, including the space that will be required to later to record location, boundary definitions, and timing information. (See *formatted capacity* for comparison.)

V

VOICE COIL - A fast and reliable actuator motor that works like a loud speaker, with the force of a magnetic coil causing a proportionate movement of the head. Voice coil actuators are more durable than their stepper counterparts, since fewer parts are subject to daily stress and wear. Voice coil technology is used in all Quantum drives.

W

WEDGE SERVO – The position on every track that contains data used by the closed loop positioning control. This information is used to fine tune the position of the read/write heads exactly over the track center.

WINCHESTER DISKS – Former code name for an early IBM hard disk model, sometimes still used to refer to hard drives in general.

WRITE ONCE – An optical disk technology that allows the drive to store and read back data, but prevents the drive from erasing information once it has been written.

X

XT – On the bus level, the original 8-bit version of what is now the AT bus.

INDEX

A

abbreviations 1-1
acoustical characteristics 4-6
actuator lock 5-4
adapter board 6-1
Adaptive Caching 5-14
Adaptive segmentation 5-14
air filtration 5-4, 5-5
AIRLOCK 2-2
AIRLOCK®5-4
alternate status register 6-13
Analog to Digital converter (A/D) 5-8
ARR (automatic read reallocation) 6-35
ARRE bit 5-23
automatic read reallocation 5-23
Average rotational latency 2-2
Average seek time 2-2
AWRE (automatic write reallocation enabled) 6-35

B

base casting 5-3
BIOS 3-14, 6-1, 6-2
BIOS requirements 3-14
bits, status register 6-17
block diagram 5-6
Buffer Controller 5-8
burst correction 5-20
bus interface connector 3-10
busy bit 6-17
Byte Interleaving 5-18

C

cable select (CS) jumper 3-5
cable select feature 6-1
Canadian Standards Association (CSA) 2-3
CE (cache enable) 6-35
check bytes 5-19
Check Power Mode (E5H) 6-25
clean room 5-1
CMOS setup 3-14
command
identify drive parameters 6-26
command block registers 6-14
command codes and parameters 6-18
command descriptions 6-19
Command Register 6-17
command register 6-17
Configuration 6-32
configuration command data field 6-34
control block registers 6-13
cooling fan 3-9
correctable double burst errors 5-20, 5-21
correctable read errors 4-8
corrected data bit 6-17
crash stops 5-4
cross checking (XC) 5-18
cylinder high register 6-15
cylinder low register 6-15

D

daisy-chain 2-2
data port register 6-14
data request bit 6-17
Data Synchronizer 5-10, 5-11
Data transfer operations 5-9
data transfer rates 4-2
Data-transfer rate 2-2
DC motor assembly 5-3
DCAIA 5-7, 5-13
DCAIA buffer controller 5-14
DCR (disable early correction) 6-35
DCSIA block diagram 5-7
defect management 5-23
defective sectors 2-2
device control register 6-13
diagnostics 6-21
Digital to Analog converter (D/A) 5-9
dimensions 4-6
DisCache parameters 6-35
disk capacity 3-14
disk errors 4-8

disk stack assembly 5-3
DOS 4-2
double burst errors 5-19, 5-20
DP8491 Read Channel Device 5-10
drive address register 6-14
drive electronics 5-6
drive mechanism 5-1
drive parameters 6-36
drive ready bit 6-17
drive seek complete bit 6-17
drive select (DS) jumper 3-5
drive write fault bit 6-17
drive/head register 6-16
dynamic offset 5-12

E

ECC 5-18
ECC algorithm 5-23
ECC Check Bytes 5-18
ECC correction span 6-36
ECC error handling 5-23
ECC interleaving 5-18
EEC (enable early correction) 6-35
EMC testing 2-3
ENDEC 5-8
environmental specifications 4-7
error bit 6-17
error detection, correction 5-17
error rates 4-8, 5-17
error recovery parameters 6-35
error register 6-14
error reporting 6-37
errors, single burst 5-19
European Standards 2-3
execute drive diagnostic 6-21

F

faceplate 3-1
Federal Communications Commission 2-3
firmware 5-12
firmware features 5-14
flex circuit. 5-4
format track 6-20
formatted capacity 4-2

H

headstack assembly 5-4
host software reset 6-13
humidity 4-7

I

IDE 3-6
IDE interface controller 5-9
IDE-bus interface 6-1
identify drive 6-26
identify drive parameters 6-26
Idle Immediate Mode (E1H) 6-24
Idle Mode (E3H) 6-25
illustrated parts breakdowm 5-2
index bit 6-17
initialize drive parameters 6-22
input power connections 3-10
INTERFACE CONTROL 5-14
interface, IDE-bus 6-1
interleave 5-18

J

jumper configurations 3-4
jumper options 3-5

L

landing zone 5-4
Logical Cylinders 3-14
Logical Heads 3-14
Logical Sectors/Track 3-14

M

maximum screw torque 3-8
media-defect mapping 2-2
microcontroller 5-7
Microcontroller Interface 5-9
microcontroller Interface 5-8
motherboard 3-11
mounting 3-7
mounting screw maximum torque 3-7
mounting screws 3-8
MTBF 4-8
multiple random errors 5-22
multiple sector reads 6-19
multiple zone recording 5-3
multiple-sector writes 6-20

N

noise 4-4
nominal conditions 4-3
Non-repeatable errors 5-23
NRZ 5-13
number of retries 6-35

P

PCB power 6-24
PE (prefetch enable) 6-35
Peak Detect 5-11
positional information 5-12
positioning information 5-12
power requirements 4-5
power sequencing 4-4
pre-amplifier 5-11
pre-compensator 5-11
precompensator circuit 5-8
preventive maintenance 4-8
PRODUCT OVERVIEW 2-1
programmed I/O 6-11
Pulse Detector 5-11
PWM 5-8

Q

Quantum configuration key 6-35

R

R/W head matrix 5-11
RC (read continuous) 6-35
RD GATE 5-10
read buffer 6-23
read channel description 5-13
read configuration 6-32
read defect list 6-30
read long 6-19
read multiple 6-22
read preamplifier 5-4
read sectors 6-19
read verify sectors 6-20
read/write ASIC 5-11
read/write operations 5-13
reallocated data errors 4-8
Recalibrate 6-19
Reed-Solomon error correcting code 2-2
register
  error 6-14
  reset 6-13
  ripple 4-4
  rotary positioner 5-4
  rotary positioner assembly 5-4
  RUEE (reallocate uncorrectable error enables) 6-36

S

Sector Count Register 6-25
sector count register 6-15
sector data field 5-18
sector number register 6-15
seek 6-21
Sequencer 5-8
servo burst and track information 5-12
servo bursts 5-12
Servo Controller 5-9
servo system 5-12
set configuration 6-33
set configuration without saving to disk 6-34
set features 6-29
set multiple mode 6-23
SETUP program 3-14
shipping container 3-2
shock 4-7
slave present (SP) jumper 3-5
Sleep Mode (E6H) 6-25
specifications 4-1
Standby Immediate Mode (E0H) 6-24
status register 6-16
  busy bit 6-17
  corrected data bit 6-17
  data request bit 6-17
  drive ready bit 6-17
  drive seek complete bit 6-17
  drive write fault bit 6-17
  error bit 6-17
  index bit 6-17
  status register bits 6-17
storage capacity 2-1
supply voltages 4-4
syndrome values 5-19
syndromes 5-23
sys tem startup 3-14

T

tampering with the HDA 5-1
Technischer Uberwachungs Verein (TUV) 2-3
temperature 4-7
theory of operation 5-1
timebase generator 5-11
timing diagram 6-8, 6-9, 6-10
timing specifications 4-3
transient voltages 4-4
TTL logic 6-2

U

uncorrectable double burst errors 5-20, 5-21
uncorrectable read errors 4-8
Underwriters Laboratory (UL) 2-3
UNIX 4-2
UNPACKING 3-2

V

ventilation requirements 3-7
vibration 4-7

W

warranty (void) 5-1
WCE (write cache enable) 6-36
WCS 5-8
write channel operation 5-13
write driver 5-4, 5-11
write long 6-20
write multiple 6-22
write sector 6-20
WriteCache®2-2

Quantum®

QUICK FAX - READER COMMENT FORM

PLEASE FAX YOUR COMMENTS AND SUGGESTIONS FOR IMPROVING THIS PUBLICATION TO:
QUANTUM TECHNICAL PUBLICATIONS (408) 894-6159

Manual Part Number	Manual Title	Current Date
81-113668-01	Quantum Cyclone 2.1/4.3/6.4AT Product Manual	

How did you use this publication?

☐ Reference ☐ Sales

☐ Installation ☐ Maintenance

What is your occupation?

Do you wish a written reply?

☐ Yes ☐ No

What is your overall rating of this publication?

☐ Very Good ☐ Fair

☐ Good ☐ Poor

Is the material presented effectively?

☐ Well organized ☐ Fully covered ☐ Clear

☐ Correct ☐ Well Illustrated

Locating the information you were looking for was?

☐ Easy ☐ Difficult

Please enter your comments and suggestions for improvement below. If you were in any way dissatisfied with this publication, we would like to know why. Be specific, if possible; give page, column, and paragraph number references where applicable.

.....

.....

.....

.....

.....

.....

.....

.....

.....

All comments and suggestions become the property of Quantum Corporation.

Your name and FAX NUMBER:

.....

.....

.....

Thank you for your interest. We will consider all comments and will incorporate changes when appropriate.